AMD RYZEN 7 2700X
Reviewed: 2nd-gen Zen moves AMD to the top PG. 74

VIRTUALIZE YOUR PC
Why you should be going virtual more often PG. 36

SPECTRE & MELTDOWN
Find the performance hit for your machine PG. 62

MAXIMUM PC

MINIMUM BS • JUNE 2018 • www.maximumpc.com

OPTIMIZE YOUR UPGRADES

Smash bottlenecks to unleash your PC's full potential

- Find & fix slowdowns
- CPUs, GPUs, & RAM
- Benchmark like a pro

STEP-BY-STEP GUIDE
Build a 2nd-gen Ryzen system right now! PG. 68

REAL-TIME RAY TRACING
What the latest tech means for the future of gaming PG. 46

FUTURE

High Performance Gaming PC

Never say die.
Go your own way.
Dodge responsibilities.
Choose freedom.
Lose the police.
Escape from reality.
Live in a fantasy.
Build landmarks.
Brave monsters.
Rescue friends.
Turn them into enemies.
Come back from the dead.
Save the world from ruin.

ONE LIFE
LIVE IT LARGE

CORSAIR

The Latest Evolution of the Most Customizable and Advanced Case in the World

MILLENNIUM

Aluminum Chassis with Steel Frame

Aluminum or Tempered Glass Side Panels

Toolless Upgradeability

PC GAMER EDITOR'S CHOICE ★★★★★

KICK ASS!

intel CORE i7 8th Gen

Powered by Intel® Core™ i7 Processors

FREE Lifetime 24/7 Service & Support
Online Financing Available

Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.

Customize Yours Today at ORIGINPC.com

Patented ORIGIN PC Variable Mounting Technology

Overclockable Laptops

with Best-In-Class Performance

EON15-X

EON17-X

Overclockable CPU and GPU

6-Core Desktop Processor

ORIGINPC.COM

@originpc

ORIGIN
HIGH-PERFORMANCE PCS

© 2017 ORIGIN PC Corporation. All rights reserved. ORIGIN "O" Symbols are trademarks of ORIGIN PC Corporation.
All other trademarks are property of their respective owners.

SUBSCRIBE TODAY!
see PG. 42

inside

JUNE 2018

24 OPTIMIZE YOUR UPGRADES

24 OPTIMIZE YOUR UPGRADES
Identify bottlenecks and break through into pro-grade benchmarking.

36 VIRTUAL BEING
Nothing to do with our attempts to clone our editor, instead our guide takes you through the ins and outs of virtual machines.

46 REAL-TIME RAY TRACING
The next big leap in graphics processing is upon us—well, almost...

QUICKSTART

12 THE NEWS
Nvidia's partnership deal under fire; Russia at war with app; more.

18 THE LIST
Maximum PC's top eight PC and gaming shows.

The HTC Vive Pro is virtually the most repairable headset we've seen.

R&D

54 AUTOPSY
We take a good look around the HTC Vive Pro—inside and out.

56 HOW TO
Protect against power loss on a Pi; check performance hit after Meltdown and Spectre fixes; more.

68 BUILD IT
Step-by-step guide to building a second-generation Ryzen rig.

LETTERS

22 DOCTOR

94 COMMENTS

IN THE LAB

76 ASUS ROG ZEPHYRUS M GM501

80 INTEL HADES CANYON NUC8I7HVK

83 INTEL SSD DC P4510 2TB

90 FINAL FANTASY XV WINDOWS EDITION

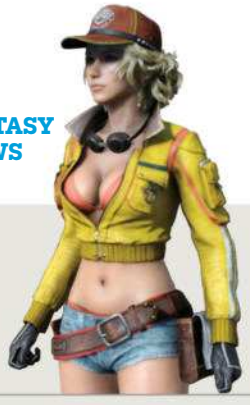

At SNHU, we engineer success. Not just inside the classroom, but outside as well. Our STEM programs feature expert faculty, high-tech labs and simulators, exciting internship opportunities, and a dedication to career development that begins on day one. SNHU's current STEM-related programs include:

Aeronautical Engineering | Air Traffic Management | Aviation Management | Aviation Operations and Management
Computer Science | Construction Management | Electrical and Computer Engineering | Mechanical Engineering

Cutting edge resources. State-of-the-art facilities. New ways to engineer success. At SNHU, we take pride in providing our students with the tools and experience they need to launch successful careers. Come see for yourself! Visit the SNHU campus today.

Southern New Hampshire University

College of Engineering, Technology, and Aeronautics

snhu.edu/stem | admission@snhu.edu | 603-645-9611

SAGER
GAMING NOTEBOOKS

SAGER RECOMMENDS WINDOWS

BLAZING PERFORMANCE

NP9877 NOTEBOOK

$2,899

FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- 8th Generation Intel® Core™ i7-8700K Processor
- 17.3" Full HD, 144Hz Matte Display with G-SYNC
- Opt. 17.3" 4K QFHD Matte Display with G-SYNC
- NVIDIA® GeForce™ GTX 1080 8GB GPU
- Opt. Dual NVIDIA® GeForce™ GTX 1080 GPUs
- 16GB DDR4-2400MHz Memory
- 256GB WD Black NVMe M.2 SSD + 1TB 7200 RPM Hard Drive
- Full Color Programmable Backlight Keyboard
- Killer Wireless-AC 1535 + Bluetooth

$2,399

NP9175 NOTEBOOK FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- 8th Generation Intel® Core™ i7-8700K Processor
- 17.3" Full HD, 144Hz Matte Display with G-SYNC
- Opt. 17.3" 4K QFHD Matte Display with G-SYNC
- NVIDIA® GeForce™ GTX 1070 8GB GPU
- Opt. NVIDIA® GeForce™ GTX 1060 or 1080 GPU
- 16GB DDR4-2400MHz Memory
- 256GB WD Black NVMe M.2 SSD + 1TB 7200RPM Hard Drive
- Full Color Programmable Backlight Keyboard
- Killer Wireless-AC 1535 + Bluetooth

$1,999

NP9155 NOTEBOOK FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- 8th Generation Intel® Core™ i7-8700K Processor
- 15.6" Full HD 144Hz Matte Display with G-SYNC
- NVIDIA® GeForce™ GTX 1060 6GB GPU
- Opt. NVIDIA® GeForce™ GTX 1070 or 1080 GPU
- 16GB DDR4-2400MHz Memory
- 250GB WD Blue M.2 SSD + 1TB 7200RPM Hard Drive
- Full Color Programmable Backlight Keyboard
- Killer Wireless-AC 1535 + Bluetooth

All Sager systems come with one year Limited Warranty & Lifetime Tech Support. Prices and specifications are subject to change without notice or obligation. Sager Corporate Offices 18005 Cortney Court, City of Industry, California 91748 ©2018 by Midern Computer, Inc. All rights reserved. Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. All company and/or product names mentioned herein are trademarks and/or registered trademarks of their respective companies. *Free UPS Ground Shipping valid to contiguous US order only.

Windows 10

intel CORE i7 8th Gen

$1,999

NP8375 NOTEBOOK FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- 8th Generation Intel® Core™ i7-8750H Processor
- 17.3" Full HD 144Hz Matte Display with G-SYNC
- NVIDIA® GeForce™ GTX 1070 8GB GPU
- Opt. GeForce™ GTX 1060 6GB GPU (NP8373)
- 16GB DDR4-2400MHz Memory
- 256GB WD Black NVMe M.2 SSD + 1TB 7200RPM Hard Drive
- Per-Key RGB Backlight Keyboard
- Intel® Dual Band Wireless-AC 9560 M.2 AC + BT
- Slim design with only 0.98 inches thin

$1,899

NP8955 NOTEBOOK FREE UPS GROUND SHIPPING*

- **Windows 10 Home**
- 8th Generation Intel® Core™ i7-8750H Processor
- 15.6" Full HD 144Hz 5ms IPS Matte Display
- NVIDIA® GeForce™ GTX 1070 8GB GPU with Max-Q
- Opt. GeForce™ GTX 1060 6GB GPU (NP8953)
- 16GB DDR4-2400MHz Memory
- 256GB WD Black NVMe M.2 SSD + 1TB 7200RPM Hard Drive
- Full Color Programmable Backlight Keyboard
- Intel® Dual Band Wireless-AC 9560 M.2 AC + BT
- Ultra-slim design with only 0.73 inches thin

$1,699

NP8851 NOTEBOOK After $150 Instant Savings!

- **Windows 10 Home**
- 8th Generation Intel® Core™ i7-8750H Processor
- 15.6" Full HD 120Hz 3ms 94% NTSC Matte Display
- NVIDIA® GeForce™ GTX 1070 8GB GPU with Max-Q
- 16GB DDR4-2400MHz Memory
- 250GB WD Blue M.2 SSD + 1TB 7200RPM Hard Drive
- Full Color Programmable Backlight Keyboard
- Intel® Dual Band Wireless-AC 9560 M.2 AC +BT
- Ultra-slim design with only 0.73 inches thin

$1,399

NP7851 NOTEBOOK After $100 Instant Savings!

- **Windows 10 Home**
- 8th Generation Intel® Core™ i7-8750H Processor
- 15.6" Full HD 144Hz 5ms IPS Matte Display
- Opt. 17.3" Full HD 144Hz Matte Display (NP7871)
- NVIDIA® GeForce™ GTX 1060 6GB GPU
- 16GB DDR4-2400MHz Memory
- 120GB WD Green M.2 SSD + 1 TB 7200RPM Hard Drive
- Full Color Programmable Backlight Keyboard
- Intel® Dual Band Wireless-AC 9560 M.2 AC + BT

CALL US TOLL FREE: 1.800.669.1624 CUSTOMIZE & PRICING GO ONLINE: SAGERNOTEBOOK.COM

Dealer/Var, government and corporate pricing are available. Please call for details.

CB CREATIVE BLOQ

> Good design is all about making other designers feel like idiots because that idea wasn't theirs.

FRANK CHIMERO

The number one destination for
graphic design news, views, and how-tos

creativebloq.com

MAXIMUMPC

EDITORIAL
Executive Editor: Alan Dexter
Deputy Editor: Zak Storey
Senior Editor: Jarred Walton
Technology Editor: Bo Moore
Contributing Editor: Chris Angelini
Contributing Writers: Alex Campbell, Alex Cox, Ian Evenden, Phil Iwaniuk, Jeremy Laird, Connor MacLeod, Nick Peers, Les Pounder, Chris Lloyd, Neil Mohr
Copy Editor: Katharine Davies
Editor Emeritus: Andrew Sanchez

ART
Art Editor: Fraser McDermott
Image Manipulation: Gary Stuckey
Photography: Future Photo Studio

BUSINESS
US Marketing & Strategic Partnerships: Stacy Gaines, stacy.gaines@futurenet.com
US Chief Revenue Officer: Luke Edson, luke.edson@futurenet.com
East Coast Account Director: Brandie Rushing, brandie.rushing@futurenet.com
East Coast Account Director: Michael Plump, michael.plump@futurenet.com
East Coast Account Director: Victoria Sanders, victoria.sanders@futurenet.com
East Coast Account Director: Melissa Planty, melissa.planty@futurenet.com
East Coast Account Director: Elizabeth Fleischman, elizabeth.fleischman@futurenet.com
West Coast Account Director: Austin Park, austin.park@futurenet.com
West Coast Account Director: Jack McAuliffe, jack.mcauliffe@futurenet.com
Director, Client Services: Tracy Lam, tracy.lam@futurenet.com
Director, Retail Sales: Bill Shewey, bill.shewey@futurenet.com

PRODUCTION
Head of Production UK & US: Mark Constance
Production Controller: Vivienne Calvert
Project Manager: Clare Scott
Production Assistant: Emily Wood

FUTURE US, INC.
11 Hanover Square, 14th Floor, New York, NY 10005, USA
www.futureus.com

Global Chief Revenue Officer: Charlie Speight
Vice President, Marketing & Operations: Rhoda Bueno
Finance Director: Ryan Lamvik
HR Generalist: Carla Marcos

SUBSCRIBER CUSTOMER SERVICE
Maximum PC Customer Care,
Future Publishing, PO Box 2024, Langhorne, PA 19047
Website: http://myfavoritemagazines.com
Tel: 844-779-2822
Email: contact@myfavoritemagazines.com

BACK ISSUES
Website: http://myfavoritemagazines.com
Tel: +44 344 848 2852

REPRINTS
Future US, Inc., 11 Hanover Square, 14th Floor, New York, NY 10005, USA
www.futureus.com

Next Issue On Sale June 26, 2018

©2018 Future US, Inc. All rights reserved. No part of this magazine may be used or reproduced without the written permission of Future US, Inc. (owner). All information provided is, as far as Future (owner) is aware, based on information correct at the time of press. Readers are advised to contact manufacturers and retailers directly with regard to products/services referred to in this magazine. We welcome reader submissions, but cannot promise that they will be published or returned to you. By submitting materials to us, you agree to give Future the royalty-free, perpetual, non-exclusive right to publish and reuse your submission in any form, in any and all media, and to use your name and other information in connection with the submission.

FUTURE
Connectors. Creators. Experience Makers.

Future plc is a public company quoted on the London Stock Exchange (symbol: FUTR)
www.futureplc.com

Chief executive **Zillah Byng-Thorne**
Non-executive chairman **Peter Allen**
Chief financial officer **Penny Ladkin-Brand**

Tel +44 (0)1225 442 244

Alan Dexter

EVERYTHING HAS CHANGED, YET STILL THE SAME

IF THERE'S ONE THING guaranteed to put you, our beloved readers, on edge, it's change. You know what you like, and change could mean that thing, whatever it is, could disappear. So the fact that it's my face staring back at you at the top of the editor's intro, and not Tuan's, will have many of you reeling about the potential implications. Let me reassure you, though: There's nothing to worry about. In fact, while this page is indeed different, you'll be hard pushed to spot any real change in the magazine.

There's a very good reason for this, and it's that I've been doing the day-to-day editing of the magazine for the last three years. Tuan was the editor-in-chief, and as such oversaw the magazine and *Maximum PC's* web presence. However, as time moved on, he was forced to focus more and more on the latter, although he still loved penning the editor's intro and replying to readers' letters.

Speaking of readers' letters, one question we get asked a lot is, "What's bottlenecking my system?" So, it seems only fair that we finally get around to giving you the definitive answer, and show you how to find out exactly where your bottlenecks are for your rig as a whole, and also how to decode the various system resources, so you can find the definitive answer for specific cases.

Once you have run a few of these tests yourself, under various different workloads, you should start to see a trend where one particular subsystem is holding your machine back. This should then inform you as to what sort of upgrade would have the biggest impact on your system overall. Whether that's the processor, graphics card, RAM, or your storage subsystem.

It seems only fitting that my first issue as editor coincides with the emergence of real-time ray tracing, because it was ray tracing that first got me seriously interested in computers. The industry has moved on incredibly since I first set my machine up to render floating metal balls above a checkered floor, and along with improved rendering techniques, we're now looking at rendering what used to take hours for each frame at 60fps. Impressive. We're just waiting on Nvidia to release some consumer-level Volta GPUs now.

This issue we've also got a feature on doing more with virtualization, a tutorial on assessing what sort of impact the Meltdown and Spectre patches have on your performance, and a review of the new AMD Ryzen 7 2700X. There's always a lot going on in the PC industry, but change is the thing that keeps it interesting and fresh. Embrace the new, and let's see where it takes us.

Alan Dexter is Maximum PC's *executive editor and a punisher of hardware. He's been a tech journalist for over 20 years, and has no problem upsetting the PC industry as a whole.*

↘ submit your questions to: comments@maximumpc.com

THE NEWS

Nvidia's Partnership Deal Under Fire

Feathers fly as AMD preaches freedom, and lays into rival

THE DISCRETE GPU market is large and lucrative, and there are really only two serious players at the high-end: AMD and Nvidia (at least until Intel also pitches in). As you might expect, there is a robust rivalry between the two companies. Nvidia has the dominant market share, currently about 70 percent. Recently, however, that rivalry has become rather heated—and public. The source of the conflict is Nvidia's new GeForce Partner Program, GPP.

The GPP is a deal between Nvidia and OEMs to collaborate on Nvidia-powered cards, working with Nvidia engineers, co-operating on marketing, launches, and so forth. Nothing unusual about that—the industry is full of such partnership deals. The company wants to corner the gamer market; specifically, it wants the gamer brands to sport GeForce hardware.

Then the story broke that this partnership had tighter restrictions than were made public. In fact, it is alleged that signing up for the GPP also meant signing up to just sell Nvidia cards under gaming brands—you couldn't sell AMD-based gaming hardware at all if you wanted to join their gang. In return, the OEMs would receive priority shipments of GPUs, discounts, and marketing resources.

This worrying accusation was leaked off the record from a number of OEMs. Because Nvidia is such a strong player in the market, being left out of the loop would be painful for any manufacturer. Publicly, Nvidia says that its GPP "isn't exclusive. Partners continue to have the ability to sell and promote products from anyone."

AMD has responded—very publicly. Its general manager of Gaming, Scott Herkelman, posted a message that talks, in rather flowery terms, of freedom of choice, openness, and innovation. He talks of freedom from "gamer taxes," and the promotion of open-source technologies. The message does not mention Nvidia, and makes no allegations. However, it rather pointedly does talk about AMD fostering real partnerships with real consistency, ones free from anti-gamer and anti-competitive strings. The post came complete with an image of a Radeon graphics card emblazoned with the logo "freedom of choice." It is strong stuff.

Meanwhile, we have an interesting move by Asus. It launched a new gamer brand of graphics cards, the AREZ (after Ares, the Greek god of war, but for some reason spelt with a "z"). These new boards are, as was quickly apparent, essentially Radeon-powered ROG boards under a new flag. If Asus is a member of the GPP, and its restrictions are as alleged, this makes sense. MSI has also removed the "gaming" tag from its Radeon cards.

All this is rather sad. Nvidia has an enviable position in the market, and really has no need at all to start playing so rough. No OEM will want to put its supply of GeForce hardware in jeopardy, so when offered the GPP to sign, it was an offer they obviously found hard to refuse. However, they wouldn't be happy about it, making a leak inevitable.

If the allegations over the fine print in the GPP deal are proved true, then this could not only prove embarrassing for Nvidia, but it could also get rather expensive. Anti-competitive practices do not go down well in court, as Intel has learned the hard way. Business is war, or so they say, but like war, it has its own rules. So far, Nvidia has kept quiet, but the story won't go away. The company's best bet may be to come clean, blame an "enthusiastic" mid-level executive, and redraft the deal before anybody thinks to lawyer up. It's only a matter of time before we all get to read the full text of the GPP deal.... **–CL**

The new Asus AREZ—it's the ROG board in all but name. The rebranding is a consequence of Nvidia's controversial GPP program.

The Genuine Article

Our legacy is simply building award-winning PCs, meticulously crafted for extreme performance. Now it's time to get yours.

- **70+** Editors' Choice Awards
- **20** Year legacy building the world's fastest custom PCs
- **560** 5-star reviews on ResellerRatings.com
- **1** Expert builder per order
- **0** Installed bloatware
- **∞** US-based phone and email support

Proudly featuring Intel® Optane™ Memory

VELOCITY micro

Configure now at VelocityMicro.com

Online financing available

intel OPTANE MEMORY

Intel, the Intel logo and Intel Optane are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.
Velocity Micro, the Velocity Micro logo, and Raptor are trademarks of Velocity Holdings, LLC.

RUSSIA AT WAR WITH APP

20 million blocked addresses and counting

THE RUSSIAN STATE is trying to shut down an encrypted messaging app called Telegram. The app started life in Russia, but moved to Amazon Web Services and Google Cloud servers outside the country after a law banning all encryption services the state can't crack was passed (ostensibly to fight terrorism).

To stop its citizens using Telegram—it has 14 million local users—the Russian regulator, Roskomnadzor (RKN), has taken the heavy-handed approach of blocking huge numbers of IP addresses, and demanding that servers remove access. The swathe of blocked addresses has caused considerable collateral damage inside Russia; sites brought down include Windows Update, Xbox Live, Twitch, Slack, Soundcloud, Spotify, Nintendo, and a myriad others. It also causes problems for Amazon and Google services themselves, and it has been reported that RKN even lost some of its own websites. RKN has denied that it is to blame for any third-party problems. The effect, so far, has been the blocking of about 30 percent of Telegram's network.

We've been here before. Zello, an app that turns your cell into a two-way radio, also fell foul of the Russians. It also jumped IP addressees in an effort to escape sanctions, before the Russians brought pressure to bear on its Amazon host. It won that round, and Zello pulled out of Russia. Telegram is much bigger, and Amazon and Google have resisted the pressure so far. How far will the Russians go? How much of the Internet will they break? More importantly, how much interference will the big hosts allow?

You can get around the ban by firing up a proxy server or VPN, something the Russians haven't addressed fully yet, although RKN has begun to issue warnings to hosts of these services, too. Meanwhile, the originator of Telegram, Pavel Durov, has promised to donate Bitcoin to those who host proxy and VPN services. **–CL**

BLACK OPS 4 HAS NO SOLO

BLAME THE SUCCESS OF PUBG

PLAYERUNKOWN'S BATTLEGROUNDS' particular take on the battle royale format has caught the imagination of gamers, and made its creators a fortune. *Fortnite* followed soon after, and has made similar amounts of cash ($100 million a month, and over). Now everybody wants a slice. There's only so much game development to go around, and one early casualty has been the single-player campaign in the upcoming *Call of Duty: Black Ops 4*. Instead, we will get a battle royale mode, and more zombies. The solo game simply can't be finished in time for the scheduled October release. It remains to be seen if we get it at all.

This changes the nature of the game, turning it into yet another multiplayer offering. The *CoD* solo campaigns have had an immersive, cinematic quality that you don't get in the frantic multiplayer mode. Other games jumping aboard the bandwagon include *Red Dead Redemption 3* and *Battlefield 5*, which will also feature battle royale modes. **–CL**

OBAMA APPEARS ON DEEPFAKE VIDEO

A warning of what machine learning can do

WE KNOW HOLLYWOOD can produce realistic digital effects, but it used to take months of work by specialists. All that is changing—apps such as FakeApp make sophisticated effects possible in a fraction of the time and with minimum effort. Witness an example put together by Jordan Peele and Jonah Peretti to highlight the dangers of deepfake videos. It features Obama appearing to say, well, whatever the creators wanted him to say. It's not as convincing as some would have you believe, but it's not bad at all. It would certainly pass many people by if inserted into a Facebook news feed. And that's the scary part.

You'll be hearing a lot more about machine learning and the neural network in the future, and one consequence will be the loss of verisimilitude for online videos. How can you be sure that what you see is real, when artifice is so easy? We will have to adjust our approach to what we view. **–CL**

Tech Triumphs and Tragedies

A monthly snapshot of what's good and bad in tech

TRIUMPHS

MORE MINI RETRO GEAR
Following the huge success of the NES and SNES Classic, Sega is launching a Mega Drive Mini.

YOU CAN OPEN THE CASE
The FTC has confirmed that "warranty void if removed" stickers are illegal, and sent out warning notices to companies.

AMAZON PRIME HITS 100M
Free delivery and constant nagging pays off. Amazon doesn't usually give out stats, but couldn't resist boasting about this.

TRAGEDIES

YOUR TV WILL SNOOP TOO
The new TV broadcast standard allows smart TVs to track what you watch and help target ads.

TESLA 3 HALTED AGAIN
Production targets are way off, with a huge waiting list. The solution? Hire more people to build the cars instead of robots.

SPECTRE/MELTDOWN COST
Benchmark figures are emerging of the performance hit patches make—mostly OK, but heavy I/O can take a 20-plus percent drop.

Jarred Walton

TECH TALK

Ryzen's First Year Was Awesome and Horrible

I'VE BEEN TINKERING with PCs since I bought my first custom 286 in 1990. For decades, the CPU was the main factor in determining how fast your PC was, and while GPUs are more important now for gaming and machine learning, I still love checking out new entrants to the CPU field.

Last year was the most excitement I can recall. Ryzen brought AMD's CPU division back to relevance, something it hadn't been since the early days of the FX-series. Ryzen was more than a 50 percent improvement in IPC (instructions per clock) over AMD's last CPU architecture. It wasn't always faster than Intel, and you need to be running multithreaded applications to fully benefit, but AMD forced Intel to finally move beyond four-core CPUs for the mainstream market. Second-gen Ryzen might even push Intel to release a mainstream eight-core processor, judging by current rumors.

But for all the good that Ryzen brought, there was a dark, less-discussed element. I'm talking about motherboards and firmware. If you weren't part of the initial wave of Ryzen 7 CPUs and motherboards, you hopefully missed most of the growing pains. RAM incompatibilities were at the top of the list, with failure to run at supposedly supported memory speeds and timings being all too common. Too bad that wasn't the only struggle I encountered. I had two motherboards completely fail on me after less than six months of use. And there's more.

Even now, powering up several of the X370 and B350 systems that have been sitting dormant is a crapshoot on whether they POST successfully or need to have the BIOS reset and memory timings re-entered. If you want to run JEDEC memory speeds of DDR4-2133, you should be fine, but if you have a nice kit of RAM, you don't want to do that. I've stopped shutting down my Ryzen PCs except when swapping hardware, as a CPU swap requires going into the BIOS regardless.

The good news is that X470 testing so far has been much better, though I haven't been able to test more than a few boards yet. It's worth noting that all three boards (Asus, Gigabyte, and MSI) use two CPU power connections: eight-pin EPS12V and four-pin ATX12V, or in the case of MSI, dual EPS12V. Those extra connections aren't just for show; fail to populate one, and you may encounter overclocking instabilities (although we didn't—Ed).

I'm not used to seeing dual CPU power connections for ostensibly mainstream platforms, and many PSUs don't come with the requisite cables. Pay attention to this area if you're building a new system around an X470 board.

The need for the added power connections comes from AMD's improved Precision Boost 2 (PB2) and Extended Frequency Range 2 (XFR2) features, along with a higher 105W TDP on the top parts. With an X470 motherboard, especially with a decent CPU cooler, you should see higher boost clocks across all workloads. But what if you put the same Ryzen 7 2700X CPU into a B350 mobo? With a BIOS update, it should work, but you'll typically see worse performance than with X470.

It goes back to the earlier designs, for both the chipsets and boards. PB2 and XFR2 communicate with the motherboard to detect appropriate power levels, and some boards, especially B350 ones, likely won't allow the new CPUs to run at maximum performance. And at least one board I tried, a Gigabyte X370 model, seemed to be going the opposite way, with the 2700X hitting temperatures of 95 C at stock under heavy workloads. That's not good.

I'd like to think all these issues will be ironed out soon, but after a year of frustrating glitches with the 300-series and first-gen Ryzen, I wouldn't bet on it. Bottom line from me: If you're thinking about buying a second-gen Ryzen CPU, grab a new X470 motherboard to go with it.

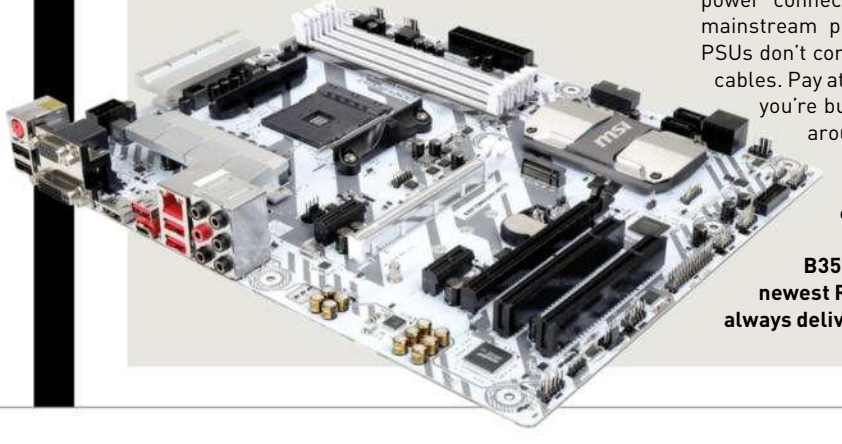

B350 mobos support the newest Ryzen CPUs but don't always deliver top performance.

Jarred Walton has been a PC and gaming enthusiast for over 30 years.

gamesradar+

One-stop destination for games, movies, TV, and more!

New Channels: Movies, TV and Cool Stuff!

www.gamesradar.com

Alex Campbell

OPEN SOURCE

How Much is Your Privacy Worth?

IT TURNS OUT, if you want to pay for an email service, web server, object storage, and a web domain, it'll cost you about $26 per month. The full picture of cost and benefit is a little more nuanced, especially since the value of convenience is hard to quantify.

There's a good chance you sat and watched Mark Zuckerberg's testimony before the House and Senate in April. If you're like me, you tried to pay as much attention as possible before you fell asleep during the four-hour Senate hearing. (I watched it on ABC's YouTube stream, and I've never seen a stream stay that stable for that long.)

If you didn't see it, you didn't miss a lot. It was b o r i n g. But if there was one takeaway, it's that Congress was finally taking the concept of online privacy seriously. (This is interesting, given that Sen. Diane Feinstein (D-CA) and Sen. Richard Burr (R-NC) introduced legislation that would have forced backdoors into crypto technology in 2016, but that's beside the point.) The point they were trying to get at is that users weren't given ample indication that their information was being used for advertising purposes. For me, and most of the tech-minded Internet, that wasn't news. The fact is the Googles, Facebooks, Twitters, and Amazons of the world house an alarming amount of our private data. In return for our data, we get services—Gmail, Instagram, Twitter bots, and cheap shipping—in return. If someone wants to opt out of social media platforms such as Facebook or Twitter, it can be challenging socially, but not impossible. Core services such as email and file hosting can be a lot more tricky.

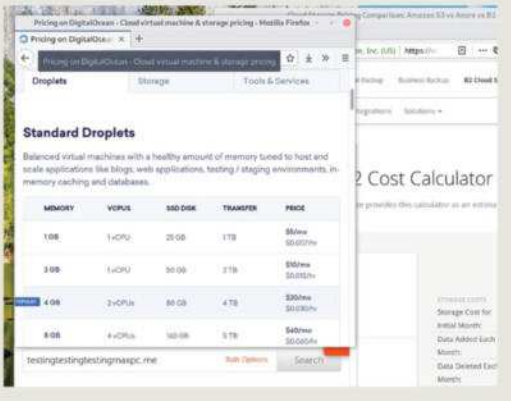

Giving up Google and social media will cost a bit.

The thing is, you can get these services without the privacy concerns, but there's a catch: You have to pay for them. I did some quick spreadsheet math to see how much it would cost to go the paid route. It turns out that the cost of privacy is surprisingly low, at just $26 per month or $332 per year.

First of all, you'll need a server. If you're already paying the big bucks for a gigabit fiber line at home, hosting at home is probably the best idea. But for those of us still with 100-plus Mb download and anemic 5Mb upload cable speeds, home hosting will render less than ideal performance. With a hosted virtual private server (VPS), you can host a range of software at a pretty low cost with high availability. With a VPS, you can host a blog with Ghost or WordPress, sync files with Syncthing, or even run a Google Drive-like application called Nextcloud.

Here's a little rundown of how I arrived at the price points. Using a Digital Ocean 2GB Droplet, VPS hosting comes to $10 per month ($120 per year). Backblaze's B2 (a competitor to Amazon S3) pricing is based on storage as well as uploads, downloads, and deletes. An initial 2TB upload with a monthly 100GB of upload, 100GB of download, and 5GB of deleted data would set you back $10.58 the first month or $158 per year. Personal domains cost as little as $3.88 per year at Namecheap. For email, Fastmail offers private email hosting—along with a calendar, contacts, and 25GB of storage—for $5 per month or $50 per year.

That's not bad, considering what you can get for each service. Granted, if you need more or less storage, your Backblaze bill will change. If you need more server performance, your Digital Ocean bill will increase.

But there are also time costs. One of the key things services such as Gmail or Google Drive provide is convenience. I don't have to spend time administering those services, and things such as Inbox by Gmail can be really handy.

So, can you ditch Big Tech and go indy online? Sure, but you'll probably miss some convenience, and definitely have to pay. And despite what Zuck said on the Hill, I've yet to find a Facebook replacement.

Alex Campbell is a Linux geek who enjoys learning about computer security.

THE LIST

TOP PC AND GAMING SHOWS

8

CES

Big gadgets, future technology, and next-gen innovation to kick off the year, every January in Las Vegas, NV. **www.ces.tech**

7 COMPUTEX TAIPEI

Running since 1981, Taipei's Computex inevitably hosts some massive tech reveals from Asia's top manufacturers. **www.computextaipei.com.tw**

6 CEBIT

Straight after Computex in June comes its European counterpart—Cebit, in Hanover, is the largest computer show in the world. **www.cebit.de**

5

PAX

Seattle. Boston. Philadelphia. San Antonio. Melbourne, Australia. PAX, running since 2004, is a massive entity with a huge following. **www.paxsite.com**

4

GAMESCOM

Cologne's Gamescom is a huge celebration of gaming, with an awesome city-wide festival tacked on to the end. **www.gamescom-cologne.com**

3

PC GAMING SHOW AT E3

LA's E3 is packed with games of all kinds, but our friends at *PC Gamer* have made it a mecca for PC enthusiasts, too. **www.pcgamingshow.com**

2 ACMSIGGRAPH

The annual summer conference of the Special Interest Group on Computer Graphics showcases emerging graphics tech with a scientific slant. **www.siggraph.org**

1

GAME DEVELOPERS CONFERENCE (GDC)

Every March, the game development community assembles in San Francisco for panels, networking, and just a little fun, too. **www.gdconf.com**

ENERMAX
POWER. INNOVATION. DESIGN.

Scan to Learn More

ENERMAX T.B. RGB™ fans
ENERMAX T.B. RGB™ fans support the latest addressable RGB-sync function to provide an incredible RGB lighting experience
High-efficiency ceramic pump with rubber cover
With TDP up to 300W
Patented flow indicator design
Patented flow indicator allows users to visually monitor cooling loop operation and coolant level
Premium 400mm weaved tubing is perfect for mainstream system builds
ENERMAX
msi

LIQFUSION

AII-IN-ONE LIQUID CPU COOLER

@enermaxusa | www.enermaxusa.com

Available Now at amazon Fry's newegg

HEAD TO

BY JEREMY LAIRD

AMD Ryzen Mobile vs. Intel 8th Gen Mobile

With the arrival of mobile variants of AMD's Ryzen chips and laptop variants of Intel's 8th Gen Core CPUs, there's some serious competition in the portable processor market. That's good news, and with it comes a few questions. For instance, does AMD's Zen architecture work as well on the move as it does on the desktop? And who makes the best processor for a wide range of portable applications? We'll also try to make sense of how these chips fit into the idiotic array of CPUs that Intel has drawn together under that 8th Gen banner.

ROUND 1

Product Range and Branding

For some time, the weakest element of Intel's operation has involved branding, marketing, and product positioning. The underlying hardware has been fantastic. However, the problem has extended beyond the ridiculous product names. Not even the fact that Intel offers the thick end of 200 desktop and mobile CPUs is the worst aspect of all that, even if that makes it almost impossible for any normal human being to keep track of Intel's CPUs, and actually understand what they are buying.

No, it's the fact that the marketing guys force infuriating decisions—like preventing most of Intel's CPU from being overclocked, or switching off HyperThreading—that grates most. Intel's 8th Gen chips are a case in point. "8th Gen" covers at least three distinct CPU architectures and thus means nothing. The small mercy is that, for now, 8th Gen Mobile is restricted to just 10 CPU models, so is just about comprehensible. This will not, obviously, last. AMD offers just four Ryzen Mobile chips, and they almost all make sense. Only the Ryzen 3 2300U with its disabled multithreading stands out as suffering from the sort of marketing non-finesse that threatens to ruin Intel's products.

Winner:
AMD by a mile

ROUND 2

Performance

In terms of grunt, there's no competition. The range-topping beast of Intel's new 8th Generation mobile range is the Core i9-8950HK, and it simply cannot be beaten. Along with the equally new Core i7-8850H and Core i7-8750H, it's based on Intel's very first six-core mobile CPU. The Core i9 model is nominally clocked at 2.9GHz, with a maximum Turbo frequency of 4.8GHz. Much, of course, will depend on specific implementation and the cooling efficiency of a given laptop fitted with the 8950HK, but those are some fairly serious numbers.

We haven't had a chance to do a direct comparison yet with AMD's finest, the quad-core Ryzen 7 2700U. However, that APU is both lower clocked and has fewer cores and threads. The Zen architecture also typically does a little less work per clock than a modern Intel Core CPU. So, the 2700U isn't in the hunt regards outright performance. We estimate the Core i9-8950HK to be in the region of 50 percent faster. That said, at any given price point that AMD's chips operate, you may well find the performance of Ryzen mobile to be entirely compelling. Quad-core versus quad-core is a very interesting comparison. Which brings us neatly to the question of value.

Winner:
Intel by a conclusive margin

ROUND 3

Value

Picking the best value proposition when it comes to mobile processors has always been a pain. That's because laptop PCs are much more highly integrated than desktops. Yes, it is possible to buy a whitebook system and add components, including a CPU, but you're much more likely to be buying the whole laptop, full specced, and ready to go. Moreover, laptops powered by AMD and Intel chips haven't always competed directly in the same market segments, with AMD systems tending to be low-cost.

With all that in mind, where does the best value lie? In theory, the new AMD Ryzen APUs make for a great all-around package that combines strong CPU performance with far superior graphics performance compared to Intel CPU models with Intel graphics (Intel's Kaby Lake G chip have AMD graphics, but are far more expensive). In practice, there are very few laptops to choose from with the Ryzen chip. Meanwhile, there are zillions of Intel laptops, and that means more choice, more economies of scale, and, ultimately, better deals. That said, if a given Ryzen-powered portable rig just so happens to precisely match your spec requirements, it might just deliver the best bang for your buck.

Winner:
Intel, thanks to more choice

HEAD

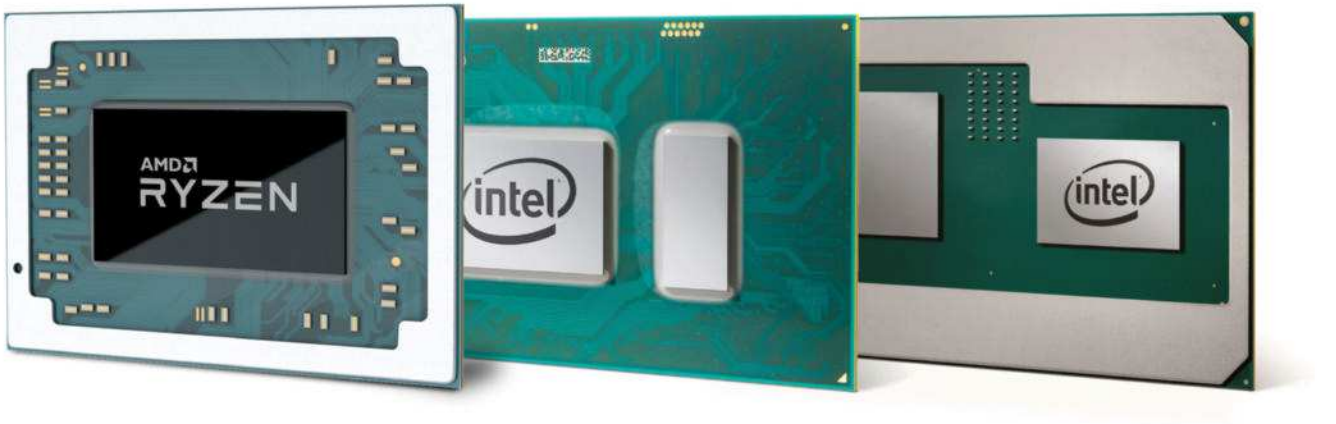

From left to right: Is AMD's Zen architecture as awesome on the move as it is on the desktop? Intel's 8th Gen mobile CPUs are lost among hundreds of other Intel models. Shock tactic: Intel added AMD graphics to its own Kaby Lake G processors.

ROUND 4

Efficiency

Of all the categories here, efficiency has to be the easiest to call. That's because AMD simply doesn't operate in the really low-power parts of the market that are dominated by Intel ultra-low voltage chips. Put simply, only Intel can provide CPUs to enable the thinnest and lightest of notebooks. The problem for this comparison is that Intel hasn't yet released 8th Generation versions of its most efficient Core M processors.

That said, there is a number of 8th Generation U Series chips from Intel that represent its second-most power-efficient offering. As it happens, those U Series Core processors line up fairly neatly against AMD's Ryzen mobile APUs. Both Intel and AMD offer dual, quad, and six-core processors, and all are rated at 15W TDP, officially. But which is more efficient? As ever, picking the question apart is awfully hard, given that laptop design plays such a major role. Handily, HP currently offers the Envy x360 laptop with both CPU options. In independent testing, the battery life performance was too close to call. For now, then, it's a dead heat. However, when Intel refreshes its Core M line with 8th Gen chips, AMD won't have anything with which it can compete.

Winner:
Intel in outright terms

ROUND 5

Graphics and Gaming

For once we have a clear winner. And it's Intel. Actually, it's AMD. Allow us to explain. Late last year, Intel rocked the PC community with an announcement involving a single processor package that combined one of its own quad-core CPUs with an AMD Radeon graphics die. The result, codenamed Kaby Lake G, is part of the 8th Gen Core range, and aimed at creating a new class of thin and light gaming laptops. It's also the fastest integrated graphics solution yet, albeit the graphics is integrated into the processor package, not the CPU itself.

Known as either a Core i7 Processor with Radeon RX Vega M GH graphics or Intel Core i5 Processor with Radeon RX Vega M GL graphics, you get either 20 of AMD's Vega graphics processing compute units or 24. That compares with a maximum of 10 Vega compute units in AMD's Ryzen mobile APUs. Admittedly, there's some debate over whether the graphics in the Intel package is truly AMD's latest Vega tech. But either way, the AMD graphics in the Intel package is clearly the more powerful proposition. Equally either way, then, it's AMD that has the most powerful graphics. Or rather Intel. You know what we mean.

Winner:
Intel via AMD. Or is it the other way around?

And the Winner Is...

Despite our reservations regarding the way it markets and positions its products, the overall result is clear. Intel wins with 3.5 points to AMD's 1.5. For the most part, that reflects Intel's much broader product offering. AMD has a single APU die that does duty for its entire range of mobile processors. It does not, for instance, offer any ultra-low voltage parts. Strictly speaking, Intel's 8th Gen Mobile range only consists of 10 models, and it too lacks models designed to go into the thinnest and lightest of portables. But it's still a broader and more capable range of mobile processors.

That gap will only grow as Intel further fills out the 8th Gen family with more and more models. On a similar note, there are simply far more Intel-powered laptops to choose from. Whatever you think of the CPU comparison, the rest of the package may well force your hand. There is, however, a caveat to all that. Where and when AMD does offer a Ryzen Mobile model, it will often be very competitive. If the spec happens to fit your needs, a Ryzen-powered portable may well be the most appealing proposition.

DOCTOR

THIS MONTH THE DOCTOR TACKLES...

> Office Location
> VR Movies
> Monochrome Desktop

Finding Office

Hi Doc. I inadvertently deleted Office 2007 Home & Student from my dated PC running Windows 10 Pro. I have the key, but no media.

I do have an older backup. However, I'd just as soon use Libre Office instead of going through the hassle of restoring the image and installing all of its updates.

Is there anywhere to download Office 2007 so I can get it back on to my computer?

–Gord MacDonald

Upcoming 4K, 144Hz displays with HDR support promise to redefine the high-end gaming experience.

THE DOCTOR RESPONDS: As of last year, Office 2007 is no longer supported by Microsoft. The company suggests upgrading to Office 365 or Office Home & Student 2016. Officially, you're out of luck. While copies of the Office 2007 disc image are floating around online, the Doc cautions against using versions found on torrent sites, many of which are infected with malware.

Incidentally, Microsoft still has a page where you can plug in a key for Office 2010 to download and install that suite (it reaches end-of-life on October 13, 2020). You can also order a replacement DVD for certain versions of Office 2010 and 2013. Moreover, copies of Office 2013 and Office 2016 are downloadable through your Microsoft account.

3D in VR?

Hello, Doc. It seems to me that VR headsets should be the ideal platform for watching 3D movies (either on Blu-ray or of one's own creation). But apparently, they can't do this.

What's so difficult about sending a left-channel output to the left eyepiece and the right-channel to the right eyepiece? **–Michael Brant**

THE DOCTOR RESPONDS: The Doc doesn't believe the challenges are entirely technical, but tied to legal complications. Either way, you compelled him to get 3D Blu-ray playback working on his test bench with an Oculus Rift.

After toying with VLC and a trial version of PowerDVD, he found success using a combination of the K-Lite Codec Pack, the latest version of LAV Filters, and madVR. All three downloads are freely available. For more information, see Tom Chapin's guide at https://bit.ly/2jA9znP. He also needed to install DVDFab Passkey for on-the-fly decryption. Without it, the Media Player Classic app kicked back an error any time the Doc tried to open his *Megamind* 3D Blu-ray.

On a computer monitor, the output appears in a split-screen mode, with one channel on top of the other. But in Guy Godin's Virtual Desktop app (available for HTC's Vive and Oculus' Rift), the combined image looks really, really cool.

Black and White

Hello, Doc. I guess a lot of things could cause the problem I'm having, but maybe you could offer a few suggestions. My system displays in full color during the POST process, and also through the course of launching Windows 10. But as soon as Windows finishes loading (the wallpaper and desktop icons are displayed) and is ready for user input, the display mode switches to black and white. In Display Settings, it is set to "Color."

Everything else works normally. I can even go into the Device Manager and see that my display adapter is problem-free with the expected driver installed.

This initially happened in the middle of a Photoshop 7.0 editing session with Internet Explorer running in background. Since then, no amount of rebooting or searching through Windows' settings has corrected the problem. In my 30-plus years of building and working

↘ submit your questions to: doctor@maximumpc.com

around PCs, I have never seen this before.

I have an Asus P5N32-E SLI Plus motherboard with an Intel Core 2 Quad Q6700 CPU, 4GB of RAM, Windows 10 Home (64-bit), an AMD Radeon R7 200-series graphics card, and an Acer G235H monitor. Before I run out and buy a new video card, I wanted get your input on what you see the issue being. **–John Hudson**

THE DOCTOR RESPONDS: Windows 10 includes a number of keyboard shortcuts for accessibility, including a Magnifier, options for changing views, High Contrast mode, Mouse Keys, Sticky Keys, and Toggle Keys. One keypress combination turns color filters on or off, and it persists through a reboot. Press the Windows logo key, Ctrl, and C together to toggle back and forth.

GPU Pricing Woes

Hi Doc. As a long-time subscriber, I always look forward to your section, and I am in a bit of a dilemma now. My custom rig includes an Intel Core i7-5930K and Nvidia's GeForce GTX 780. I have been saving to buy a 4K monitor, and I'd like a GeForce GTX 1080 Ti in order to game at its native resolution.

As cryptocurrency mining has driven up the price of high-end graphics cards, my dream of buying a 1080 Ti remains out of reach. And it doesn't look like they're getting cheaper any time soon.

Could you please help me with another option for gaming at 4K? Or should I keep waiting and saving? I have about $700 set aside, with a goal of spending around $1,200 for the monitor and card. I mostly play first-person shooters and prefer high-quality visuals over twitchy gameplay. Super-high frame rates aren't critical; playable performance is fine. **–Erik Aukerman**

THE DOCTOR RESPONDS: The Doc has some bad news and some good news, Erik. Right now, a GeForce GTX 1080 Ti is the best card for consistently playable frame rates at 3840x2160 using the highest-quality detail settings. A Titan Xp is great as well, but it's even more expensive. Once you step down to a GeForce GTX 1080 or Radeon RX Vega 64, detail settings must be compromised in certain scenarios.

Not that it matters to anyone comparing today's prices to those from a year ago. As of this writing, 1080 Ti prices start at $900, while slower Radeon RX Vega 64s land in the $800 range. It'd be difficult to recommend either option marked up the way they are.

Two things are happening, though. First, the perceived value of Ethereum is well off its highs, affecting demand for GPUs. More models seem to be in stock, even if they still command a premium. Should the trend continue, it will be a matter of time before supply catches up, nudging prices back down. Second, the introduction of Nvidia's next-generation architecture is rumored to be fast approaching. While it's too early to speculate about performance, a GeForce GTX 1080 replacement capable of smooth 4K frame rates isn't far-fetched.

There's one other reason to consider waiting for your next display subsystem: Much-delayed 4K monitors with 144Hz refresh rates and HDR G-Sync support should be shipping by the time you read this. Expect them to offer amazing visuals at gut-wrenchingly high prices.

Mobile Gaming

Dear Doc, I love *Maximum PC* for all the tech details and how-tos. I especially watch for articles on VMware, VirtualBox, DOSBox, and so on. I have all three installed for my old Windows 95/98 and DOS games. But I'd like to put old CDs on my hard drive and play those games on the road. My laptop doesn't have a built-in optical disc drive and schlepping around an external drive with a binder full of CDs isn't a solution. I would really appreciate some suggestions

A closed-loop liquid cooler should be capable of taming a stock FX-8370 CPU's 125W TDP.

for all of my virtualized operating environments, to make loading and switching between games easier. Trying to remember which platform requires what commands to load a game is frustrating.

I need a solution that does not require an Internet connection during gameplay, as I intend to use my laptop in locations where online access is either non-existent, unsafe, or outrageously priced.

– D. J. Foreman, PhD

THE DOCTOR RESPONDS: Back in his college days, the Doc had a bad habit of misplacing game discs. After losing at least two copies of *Command & Conquer: Generals*, he started archiving disc images in .iso format. Mounting these using Daemon Tools or Alcohol 120% made it possible to play his favorite titles without hunting down physical media. Perhaps the same approach would work in your virtualized environments.

Hot Chips

Hi, Doc. A few years ago, I built a PC with an AMD FX-8370, Sapphire's Radeon R9 390X, a 620W Antec HCG PSU, and an Asus M5A99FX Pro R2.0 shoved into a Thermaltake Urban S31 case. I recently started playing games more seriously, so I purchased an NZXT Kraken X31 120mm AIO cooler to alleviate the extra thermal load. I have a 120mm fan drawing air in through the radiator through the rear, another 120mm intake fan in the front, and a 240mm output fan on top. However, I am still seeing temps of up to 80 C under load, as recorded by HWiNFO. Is that normal for an FX-8370 at stock?

–Liam Andrews

THE DOCTOR RESPONDS: Temperatures that high on an FX CPU are highly unusual indeed. To start, are you certain those readings are correct? Try installing AMD's own OverDrive utility, keeping an eye on clock rates and thermal margin readings. If the processor really is getting as hot as 80 C, the available thermal margin will be consumed (plus some), and you should see the chip throttle itself to avoid damage.

In that case, it's entirely possible that your closed-loop cooler isn't installed correctly. Check its thermal paste, the mounting hardware, and any attached fans to make sure they're spinning. Go into the motherboard's BIOS and make sure voltages are set correctly.

Readers' Help

Doc, on page 22 of your Spring 2018 issue, a reader mentions an old boxed HP xw4200 Workstation, where his "32-bit Windows 10 Pro installation disc got annoyed," but doesn't provide any other description of the problem.

I have enough gray in my hair to have worked with some of those old systems, and one clue is when he mentioned having five 160GB HDDs.

Could it be that his boot order is set to a RAID controller first, rather than SATA or IDE? Back in the day, as Windows was installing, there was a function key you could press to load a RAID controller driver. I wonder if this might help Kieran Al with his issue? **–Erik Goldoff**

THE DOCTOR RESPONDS: Excellent suggestion, Erik. That specific HP workstation model included several hardware and software RAID options covering IDE, SATA, and SCSI interfaces. Perhaps Kieran was having an issue with a BIOS setting. Or maybe there is no Windows 10 driver for the storage controller those drives are attached to.

OPTIMIZE YOUR UPGRADES

Identify bottlenecks and break through into pro-grade benchmarking

By Zak Storey

COMPUTERS ARE funny things: a complicated collection of components, each working in conjunction with the rest to produce on-demand output. But as with any team effort, there is always a weak link in the chain.

As more and more people join the PC enthusiast community, it's easy to get swept up in the hype, and upgrade a very basic system with high-end componentry. And while that's not always a bad thing, it can lead to those weak links being stretched to breaking point, as top-end hardware makes its way into a system that wasn't designed to host it.

This can lead to bottlenecking, when one piece of hardware doesn't perform as well as it should, due to another component in the system being ill-equipped to handle the potential performance levels. As with the neck of a bottle, only so much can come through at once, so the power that's ready to be tapped is unused. It's usually easy to spot the part that is underperforming—a GPU, CPU, or RAID setup—but identifying which component is causing the bottleneck can be more difficult.

So, which are the most common problem parts? How do you identify them? And how do you break through that barrier? Read on to find out....

RADEON
CORSAIR
msi
SAMSUNG
VIPER
VIPER
intel

Integral System Parts

The most common cause of bottlenecking is your platform and processor

CPU

Bottlenecking is typically down to your processor. Your CPU processes everything from rendered in-game frames to memory storage requests, complex program processes, and more. It's an intricate part of the puzzle, and the easiest to become overwhelmed in any system.

Most common bottlenecks, especially in games, are due to the processor not being fast enough to handle all the data and processes thrown at it at any given time.

Six or seven years ago, most games required two to four high-powered cores to get the job done. Single-core was king, and the more instructions per clock your CPU could process, the better. Intel's well-founded single-core legacy comes from its early Nehalem architecture, which has seen no fewer than seven iterations since then, each capitalizing on improved single-core performance and support for newer hardware. Even, today Intel still holds that reputation, and despite the fact that Ryzen is now much closer to Intel's single-core performance, the gaming community still leans passionately toward Intel rather than its Team Red rival.

CPU SYMPTOMS

The most common and obvious sign of a CPU bottleneck is low frame rates in games, even when using a powerful GPU. In short, your CPU is in control of the entire game. From a simplified point of view, it runs the process logic, handles your input to the game, runs the physics for that event, then calculates the AI's response, mixes the relevant audio given the triggers activated previously, and finally makes a request to the GPU to render the frame. At which point, it returns straight back to the process logic, and starts all over again. It does all this while running the OS and background apps, and moving data in and out of various frame buffers, DRAM, cache, and off the storage medium, back and forth between the GPU. To achieve a smooth 60fps, your processor has to manage all of that 60 times per second. In more complex titles, there are even more steps added to the game logic process chain, and as multiple events occur at the same time, it's easy to see how a low-end CPU struggles to handle a modern title.

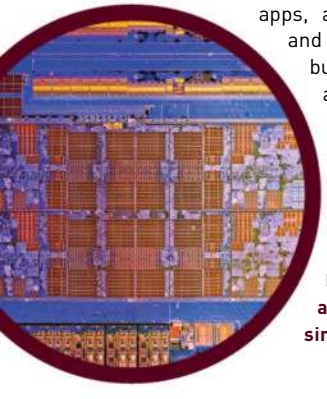

More cores don't always beat higher single-core grunt.

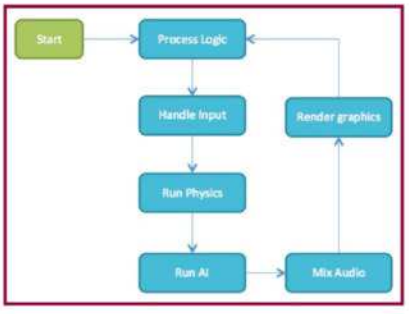

A simplified version of what your CPU goes through to produce a frame.

GAME CODING

How the processor performs in game, however, is dictated by the game devs. A CPU always operates at the same speed and frequency, processing the same number of instructions per clock every single second. If you send it a complex process to compute, filled with unnecessary jargon and redundant code, it takes up more of your processor's resources, so it takes longer to go through that chain we mentioned earlier, which in turn slows down frame rates. The developers can also allocate which processes take advantage of which cores, and how many of those cores.

SOLVING THE CONUNDRUM

How do you spot a CPU bottleneck? Typical symptoms include long in-game load times, slow texture loading, and lower frame rates than you'd expect, given the GPU you're using. On the flip side, you may notice seemingly high frame rates, but textures don't appear to render correctly, with horrendous levels of stutter and lockups.

To confirm this diagnosis, we need to do some testing and data collection. This requires a very in-depth benchmark, such as the one integrated into *Assassin's Creed Origins* or *Tom Clancy's Ghost Recon: Wildlands,* and a third-party system monitoring app, such as HWInfo or HWMonitor. *Assassin's Creed* features a benchmark HUD, which actively shows you how many cores are being utilized, and what percentage of load they're under, while also displaying you the maximum GPU load. Regardless of what you use, we're looking for the same info: CPU and GPU load.

If, for instance, you have a processor with two cores and multithreading—an old Intel Core i3, say—and all four of those logical cores are registering 100 percent load for the duration of your benchmark, with the GPU hovering around 60 to 80 percent, we can safely say the CPU is the bottleneck, because in an ideal world, the reverse would be true. This is what we call an absolute bottleneck, as you'll start to see in-game degradation, textures not loading, slow load times, and frame stuttering.

That said, you don't necessarily need all of your CPU cores running at 100 percent load to find yourself with a bottleneck; minor ones can occur due to slower single-core performance, as you'll see later with our benchmark results.

The differences are most notable at higher refresh rates, with higher powered GPUs, but they still exist at the low end. And if you play competitive FPS games, for instance, at 144Hz and above, at 1080p, single-core performance is still very much king in contrast to multicore brute force.

CAN YOU STILL BOTTLENECK?

Despite what most folk think, we've come a hell of a long way over the last seven years. And although Intel's 10 percent incremental performance increases generation upon generation haven't exactly wooed anyone, almost a decade of them has led to single-core performance increasing by a staggering 93 percent since Sandy Bridge. Combine that with access to more cores as the market heats up, thanks to the battle between Teams Red and Blue, and what you're left with is quad-core $100 i3s pumping out nearly twice as much grunt as their siblings back in 2010.

Game development has stagnated in comparison. Although graphics in the big story-driven AAA titles have improved in leaps and bounds, few developers take advantage of more than four cores. This is purely down to the fact that the majority of the market still sits on quad-core parts. What that means is that in today's world, even that $100 Core i3-8100 will be more than enough to play the vast majority of modern AAA games without bottlenecking any GPU, certainly at 4K.

Thanks for the Memory

How much is enough? What about frequency?

Memory is the support class of PCs.

Memory is an integral part of any system, and it interacts with the CPU in a very interesting way. It's essentially a hyper-fast, volatile storage medium, designed to hold data that needs to be accessed more quickly than that found on non-volatile hard drives, yet using files that are larger than those that can be held locally on the internal CPU cache.

When it comes to processors, outside of rendering tasks or simulations that operate off large data sets, it's difficult for your memory to bottleneck your system. Most everyday tasks don't process or shift enough data fast enough for the bus between the DDR and the CPU to ever become truly saturated, or for the memory to be under 100 percent load all the time.

You can encounter bottlenecks when video rendering, as large amounts of data and imported video clips are transferred from the non-volatile storage to the DDR, then the CPU and GPU, but aside from that, memory speed (frequency, in this case), for Intel at least, matters very little.

INFINITY FABRIC

AMD's Ryzen architecture, on the other hand, does take advantage of increased memory speeds for flat single and multicore performance increases, because the interweaving fabric that connects each of Ryzen's core complexes together runs at the same frequency as the memory itself. This means that the higher the memory speed, the faster the Infinity Fabric runs, and therefore the faster each core complex can transfer data and instructions between the other complexes and cores in turn. That said, there are diminishing returns from increasing memory frequency with Ryzen parts—typically, anything above 3,200MT/s on the Ryzen 1000 series will fail to see any further improvement, while the Ryzen 2000 series increases that limit to 3,600MT/s.

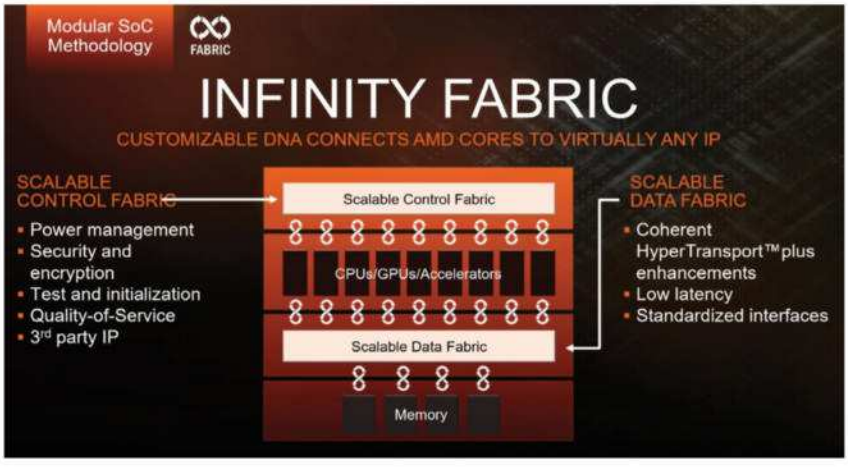

AMD's Infinity Fabric thrives on higher frequency memory.

DUAL-CHANNEL VS. QUAD-CHANNEL

Frequency? Bandwidth? Dual and quad-channel? What's going on? In reality, maximum performance is a combination of four factors: the base clock frequency at which the memory operates, the number of data transfers possible per clock (DDR means double data rate), the width of each channel or memory bus (64 bits), and the number of channels used to transfer data to and from the CPU.

Say we have a 16GB kit (2x 8GB) of dual-channel DDR4, with a base frequency of 1,600MHz. We can do the following calculation to figure out how much memory bandwidth we have. We take the 1,600 (MHz), break that down into Hz (1,600,000,000), and multiply it by the data rate (2) to get 3,200,000,000, aka 3,200MT/s (or 3,200MHz if you're in marketing). Multiply that by 64 (the width of each channel in bits), and multiply that by 2 (the number of channels), and we're left with the final figure in bits per second. Divide that by 8 for bytes, and it's 51.2GB/s, which is the maximum amount of data that can be transferred to and from your memory modules at any time. For quad-channel, you double that figure.

The important thing to note here is that no matter how many sticks of memory you have installed in your system, the CPU can only write or read from however many channels it supports at one time, so for mainstream platforms, that's two. You could have four sticks installed in a dual-channel system just fine, but all you're essentially doing is increasing the maximum memory capacity of your rig, enabling you to store more data in your DRAM—it's not faster than just running two.

As far as video editing is concerned, the more bandwidth you have the better, especially if you're working on 4K video and higher, because those files and frames are often stored directly in DRAM, and it's very easy to bottleneck a system that way.

Graphics Bottlenecks

The wonderful world of VRAM, and how a GPU does what it does

Memory is also a vital component when it comes to rendering graphics, both with integrated and dedicated GPUs. It all comes down to the notion of a frame buffer. In short, the GPU cores found in both graphics cards and CPUs rely on the frame buffer to store rendered textures, which can be called upon, when required, for quick, smooth, and efficient rendering, and the piecing together of frames in game. For a standalone GPU, the card itself features its own dedicated memory, designed specifically with that one task in mind. These memory chips typically run faster than their DRAM counterparts, and have a wider bus, which allows for a larger amount of data to pass through at any given time.

For instance, an Nvidia GTX 1080 has a 256-bit bus, with its GDDR5X chips running at a base clock frequency of 5GHz (or 10GT/s), and can transfer up to 320GB/s of data to and from the memory found on the graphics card itself—and that's all off a single channel (dual-channel is rumored to be coming with GDDR6).

Integrated graphics run off a very similar system, with both Intel and AMD's latest APUs featuring far smaller capacity memory located on the CPU's integrated GPU, and it acts in exactly the same manner.

The big problem for GPUs, however, is when the total dedicated VRAM storage capacity becomes saturated. It's not the width of the bus that's the problem, but more the total capacity not being big enough to hold all the high-definition textures and pre-rendered frames necessary to smoothly render the frames in the game you're playing. In this situation, the GPU offloads any remaining texture files on to traditional DRAM, which is far slower than its VRAM counterpart. These files have to be transferred across from the GPU to the CPU, and then to the DRAM, before going back again when they're needed. This increases CPU utilization, bumps up texture load times, and increases frame latency as the data has to go back to the GPU via the CPU. This is where you begin to get your memory bottleneck, and why, on occasion, you'll find GPUs with cores capable of driving games at two separate resolutions, often coming with multiple VRAM configurations.

LOAD RATIOS AND 4K

When it comes to GPU core performance, resolution and frames per second play a big part in overall system load. Generally speaking, the higher the frame rate, the harder the CPU has to work to process all that information effectively, and the easier it is for your CPU to become the bottleneck in a system. As the resolution increases from 1080p to 4K, that load is then offset to the GPU, and the graphics core spends more time rendering each scene (due to that 300 percent increase in the number of pixels). Because of that, the load shifts away from the CPU and on to the GPU, and simultaneously frame rates fall, because the GPU simply can't keep up.

This is important to understand, because a GTX 1080 Ti and Core i3-8100 combo is more likely to work better at 4K than at 1080p, so far as bottlenecks are concerned.

VRAM operates at a far faster frequency, with a wider bus, than traditional memory.

OVERCOMPENSATING

Having read the above, you should now have some understanding of just how important it is to buy the right GPU for the resolution you want to take advantage of, and the games you want to play. Memory and GPU cores are typically well paired. A GTX 1060, for example, combined with 3GB of VRAM, is more than enough for 80 percent of most modern AAA 1080p titles, with 6GB of VRAM being enough to future-proof you for the next four years, or just enough to scrape by on for 1440p.

The big problem, however, lies with the folk who drop the cash on a GTX 1080 Ti, only to then game at 1080p. Unless you're absolutely obsessed with playing first-person shooters at 200Hz and above, with super-low latencies, such a purchase makes little sense from a consumer standpoint. Even if you roll with the notion of it being part of a smart future-proofing decision, the reality is that HD textures are highly unlikely to reach 11GB in size any time soon, and even if they do, by that point 1080p will be positively antiquated technology in comparison to both the GPUs and the resolutions of the future.

4K shifts most of the load away from the CPU and on to the GPU instead.

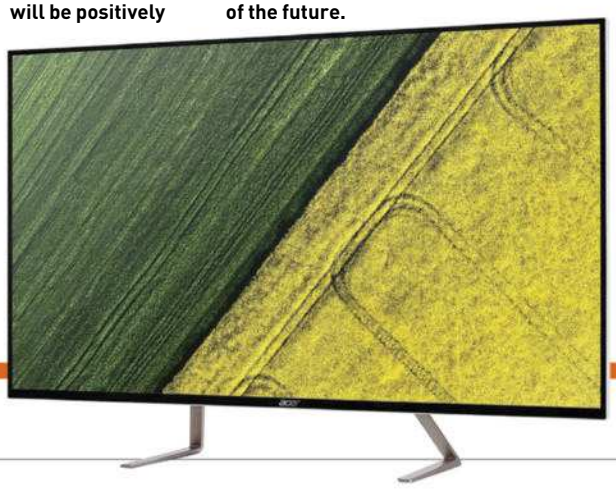

Storage Chokes

Your choice of platform dictates whether or not you can run a RAID

Storage solutions are the next biggy on our list of bottleneck conundrums, and the issue pretty much boils down to the introduction of the latest PCIe SSDs, from the likes of Samsung, causing problems for those interested in setting up multiple drives in RAID 0. Traditionally, when you wanted more speed from your system, the obvious way to do this was to take two hard drives—spinning platters or flash SSD—and set them into a redundant array of independent disks with striping. You'd take two drives, and force the data saved to be split in half and written to both simultaneously. Then you could read the original data from both drives at the same time, effectively doubling performance, in exchange for a little more vulnerability, because if one drive fails, half your data is gone.

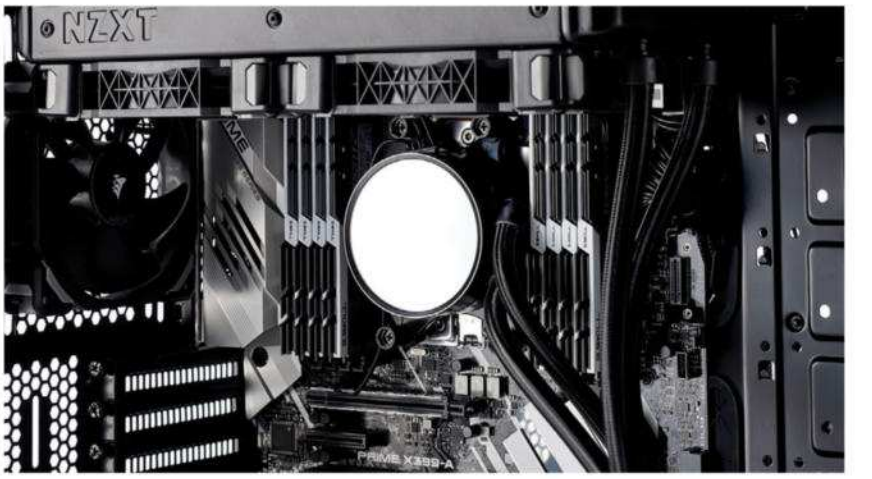

Threadripper's minimum BS attitude to RAID 0 PCIe SSDs is something we can get behind.

PCIE SSD RAID 0?

The big question is: Can you do that with PCIe SSDs? The answer: Technically, yes you can, but you won't gain any performance out of it. At least, not from most mainstream platforms (check out our results on page 34). The problem derives from how both Intel and AMD have configured the use of their chipsets and storage I/O, and it's all to do with the PCIe lane interconnects. Let's break it down.

Intel's mainstream platform takes advantage of what's known as a DMI (direct media interface). Briefly speaking, this is a direct connection between the CPU, the PCH/chipset, and the I/O storage devices in your system (SATA, USB, and M.2 SSDs). From Z170 onward, Intel has been using DMI 3.0, which is essentially four PCIe 3.0 lanes. What that means is that the maximum throughput from the chipset to the CPU via the DMI is currently 3.93GB/s (or 3,930MB/s).

All of a sudden, it's easy to see what's causing your bottlenecks when running twin M.2 PCIe SSDs in RAID 0, because there's simply not enough bandwidth to allow for a straight doubling of performance, as was possible with SSDs and HDDs through the SATA interface and DMI 2.0 (maximum throughput about 2GB/s).

AMD's Ryzen architecture runs things differently. The PCH to CPU interconnect actually only runs four PCIe 2.0 lanes, similar to Intel's original DMI 2.0, as opposed to DMI 3.0's four PCIe 3.0 lanes we see today. However, that's dedicated to onboard USB 3.1/0 and USB 2.0 devices. This is because although both Intel and AMD's processors reserve 16 PCIe 3.0 lanes directly for graphics solutions, AMD also reserves an additional four PCIe 3.0 lanes directly to the CPU for PCIe SSDs and SATA drives, with no interim chipset adding latency to the affair.

That said, it's still only four PCIe 3.0 lanes, and a 3.93GB/s max transfer rate between the CPU and its storage, meaning again, there's just not enough bandwidth.

WHAT ABOUT THE HIGH END?

At the high end, both platforms change things up quite dramatically. As both teams field processors with more and more dedicated onboard CPU PCIe 3.0 lanes (Intel's top-end chip offers a fruity 44, while Threadripper's entire range throws up 64 in total), it is indeed possible to run PCIe SSDs directly off the CPU, as opposed to running them through the chipset and DMI.

Despite Intel having multiple PCIe lanes on its chips, it hasn't supported RAID on PCIe SSDs until very recently. With Skylake-X, it introduced VROC (Virtual Raid on CPU) for the first time, allowing PCIe SSDs utilizing those direct CPU PCIe lanes to theoretically form a RAID directly on the CPU itself. However, there's a few caveats to this. Firstly, you need a VROC key—this is a physical dongle that you plug in to the motherboard to enable the feature. It's not possible to buy these separately from an official retailer, because they're typically only bundled with specific motherboards and SSDs (although you can buy one from a third party for roughly $100-150). And once that hoop's been jumped through, you can also only use Intel SSDs for VROC, which limits you to Intel's Optane 760P SSDs, which are hardly mind-blowing when it comes to storage.

On the flip side, AMD's Threadripper series packs a whopping 64 PCIe lanes directly on the chip itself, allows you to RAID whatever you like directly on the CPU without any proprietary nonsense, and is realistically the most sensible platform, where you'll see sequential speeds bounce up because of it.

RAID 0 support is limited, depending on the platform.

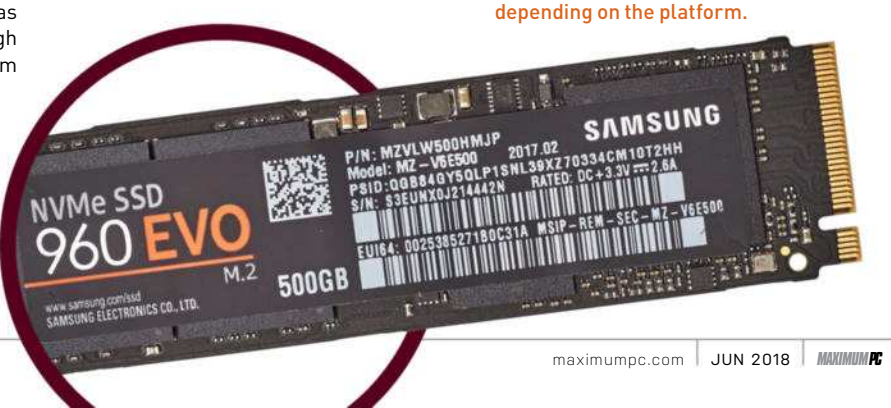

LET THE AMAZON ECHO CHANGE YOUR LIFE TODAY

Discover everything you need to know to get the most from Amazon's voice-activated speaker and its built-in digital assistant, Alexa

ON SALE NOW

FUTURE

Ordering is easy. Go online at:

www.myfavouritemagazines.com

Or get it from selected supermarkets, newsagents or bookstores

Testing Time

Let's fire up some software, and perhaps a spreadsheet or two, and figure out where exactly your bottlenecks are

CPU

There are a few obvious bottlenecks we're going to ignore, as they're processor limitations. For instance, video editing benefits from more cores and higher single-core/IPC. But the biggy is in game. Go to www.cpuid.com and download HWMonitor (dodge the Pro version), run that, and minimize everything aside from the "Utilization" section for your processor and graphics card. Here you'll see all your threads, and their current, minimum, and maximum utilization, along with your GPU utilization. Load up the gaming benchmark of your choice (Unigine's Superposition is a solid alternative at https://benchmark.unigine.com), then Alt-Tab back to HWMonitor, select the "View" drop-down, hit "Clear Min/Max," then go back to your benchmark, and hit "Run."

You're looking for whether your CPU load is at 100 percent across all threads while your GPU load is at about 60–70 percent. If it is, you'll want to look at upgrading that CPU to something with a little more grunt. We recommend a minimum of an Intel Core i3-8100, or a Ryzen 3 1300X.

HWMonitor is a vital tool for any PC junkie.

In-game fps overlays make figuring out performance woes far easier.

GPU

Again, the obvious bottleneck with the GPU comes from a lack of core performance for the resolution you want to play at. You want to hit around 40fps in game to have something of an enjoyable experience; 60fps is the sweet spot, but it's not the end of the world if some titles don't reach that—it's not suddenly unplayable. There's a ton of ways to see frame rates in game; some titles support it as an option, but if that's not possible, you can download a trial of Fraps (www.fraps.com), which includes a visual fps overlay to allow you to see your frame rate in the corner of the screen.

We can test whether a lack of VRAM is bottlenecking your system by using HWMonitor and the game you're having difficulty with. Keep HWMonitor open and those utilization windows up, and play the game you're having difficulty with. If you notice the memory is maxing out at 100 percent, either reduce the resolution of the textures, or opt for a newer GPU.

SSD

And finally, we have the wonderful world of solid-state bottlenecks. Unfortunately, as we've already discussed in depth, anyone who's running more than one PCIe SSD in something like a RAID 0 array—unless you're running those devices in a Threadripper motherboard, and have gone about extensively setting up the array for them—it's an absolute truth that you're bottlenecking both drives, while simultaneously increasing your risk of data loss. CrystalDiskMark is our go-to benchmarking software for this challenge, so head to https://crystalmark.info to download the latest version of our favorite SSD benchmark, then it's simply a case of running the QD32 Sequential speed test on the drive or array in question to identify just how fast it's running.

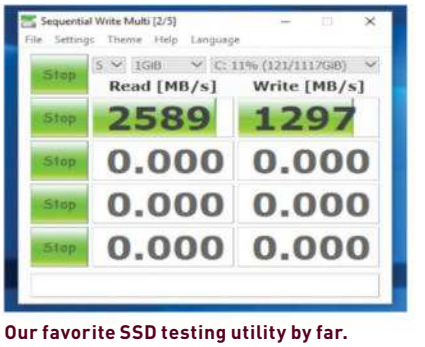

Our favorite SSD testing utility by far.

If you're running multiple PCIe drives on any mainstream board, it may also be worth checking to see whether all your M.2 slots are running in PCIe mode, and not SATA—after all, you don't want a 3,000MB/s drive to be capped at 550MB/s because of the SATA interface.

techradar.

The home of technology

techradar.com

Case Studies

Let's look at a few rigs we've specced together with obvious flaws

THE ANCIENT GAMER

INGREDIENTS

CPU	Intel Core i3-3240
Cooling	Stock
Memory	8GB (2x 4GB) DDR3 @ 1,600MT/s
GPU	2x GTX 970
Motherboard	Gigabyte GA-Z77X-UP5 TH
SSD	120GB Samsung 840 Evo
HDD	1TB Western Digital 7200rpm
PSU	550W Corsair TX550M

There are two major problems here, the first of which is the dual GTX 970s running in SLI. This forces the GPUs to operate in PCIe x8 mode, as opposed to a full-fat x16. Couple that with the fact you've got just a dual-core processor, and most games are immediately going to throttle even at 1080p.

The GPU solution is fine on its own for 1080p, but the problem lies with that processor. It would make sense to upgrade the motherboard and CPU to something with more cores and single-core grunt.

The second and possibly more alarming problem is the 550W bronze-rated power supply. By default, you should be drawing around 450W from the wall—even with this dual SLI setup—but only having 100W of headroom doesn't fill us with confidence if things spike under extreme load. Worst case scenario: The power supply can pop, and take parts with it. Best case: It just won't power on under those circumstances.

THE EAGER PROFESSIONAL

INGREDIENTS

CPU	AMD Ryzen 7 1700 @ 3.8GHz
Cooling	Stock
Memory	32GB (2x 16GB) DDR4 @ 2,666MT/s
GPU	Nvidia GeForce GTX 1080
Motherboard	MSI X370 Gaming Pro Carbon AC
SSD	1x 512GB Samsung 970 Pro
HDD	1x 1TB Samsung 960 Evo
PSU	700W Be Quiet! Pure Power 10

In this scenario, the end user has taken a single Ryzen 7 1700 and ramped it up to 3.8GHz across all 8 cores and 16 threads. Which is fine—the 1700 is more than capable of handling those kind of voltages and frequencies. However, the biggest problem here is something we haven't covered so far, and that's thermal throttling, which in its own way is a very specific form of bottlenecking.

In short, as the temperature gets too high, if your processor is stable, it immediately begins to downclock itself, to ensure it doesn't damage any of the silicon at its heart. You may think it's running at 3.8GHz, and it may very well be at startup, but as soon as you get in game, or start running some computational tasks, that frequency is going to drop, just to keep the chip safe.

On top of that, this rig has multiple PCIe SSDs installed. The Samsung 970 Pro for the OS is directly installed on to the motherboard, and will operate just fine utilizing those four dedicated lanes we spoke about earlier. The problem is the 1TB 960 Evo—although some motherboard manufacturers do allow you to run an additional PCIe SSD in an M.2 slot, they are usually forced to utilize the majority of the interconnecting lanes between the PCH and the CPU. In Ryzen's case, that's four PCIe 2.0 lanes, meaning maximum transfer rates of 2.0 GB/s. Not good.

THE EPEEN STREAMER

INGREDIENTS

CPU	Intel Core i9-7900X @ 4.7GHz
Cooling	NZXT Kraken X72 360mm AIO
Memory	8GB (2x 4GB) DDR3 @ 1,600MT/s
GPU	Nvidia Titan Xp
Motherboard	Asus Prime X299 Deluxe
SSD	1x 250GB Samsung 960 Evo
HDD	1x 1TB Crucial BX300 2.5-inch
PSU	1,200W Corsair AX1200i

In our final selection, we have a combo designed to make any tech journalist flinch. The issue here isn't that obvious at first glance, but it all comes down to the motherboard and CPU combo, combined with that hefty overclock. Although the Core i9-7900X is aptly cooled by the 360mm AIO to an impressive 4.7GHz overclock, the problem lies with the VRMs on the motherboard. There's too much voltage being pushed through them, and not enough heatsink to provide them with sufficient cooling. In short, the mobo's BIOS is forced to reduce the voltage to keep temperatures within acceptable parameters.

Although the CPU may well register at 4.8GHz, and be operating at 100 percent utilization, the actual performance can be even lower than stock in some cases. The best way to spot this is by installing HWInfo, and analyzing your VRM temperatures—anything above 110 C when under load is bad news. Run Cinebench R15 at stock and then at your overclock settings, and if you're noticing a drop in score without a drop in utilization or registered clock speed, you know it's the VRMs causing the issue.

There are two solutions: You can either drop back on the voltage, or opt for a beefier cooling solution, one that directly cools the VRMs. EKWB and some other aftermarket liquid-cooling companies offer full-cover monoblocks for both the VRMs and the CPU on most mainstream motherboards.

Bottleneck Results

Exactly how do different setups affect your performance?

HOW WE TESTED

Obviously, we no longer have a lot of the old hardware we've featured below, so much of this is synthetic, rough benchmarking, utilizing a Core i7-8700K. The theory is fairly simple: We found our old Cinebench R15 single-core figures for everything back to Sandy Bridge, and declocked the Core i7-8700K to the point at which we achieved similar, if not identical, scores to the old processors. At that point, we disabled a variety of cores and/or Hyper-Threading, depending on the configuration of the part we were trying to emulate at the time.

Other than that, our test rigs included 16GB of DDR4 for our mainstream platforms, operating at 3,200MT/s, and 32GB of DDR4 for our X399 and X299 RAID testing. Our PCIe SSDs of choice consisted of dual Samsung 960 Pro 512GB SSDs, operating through onboard M.2 slots, or Asus's included ROG Zenith adapter. As for mobos, our Intel games and RAID testing featured MSI's Z370 Godlike Gaming. For Ryzen RAID testing, we plumbed in our recent MSI X470 Gaming M7 AC. For our Intel RAID testing, we took advantage of the Asus Prime X299 Deluxe, and for Threadripper RAID testing, we grabbed our trusty Asus X399 ROG Zenith Extreme, in combination with the included PCIe SSD adapter.

All games were tested at the Ultra preset, with HD textures. 3DMark's Sky Diver was run at default for 1080p, and then adjusted to 4K for the 4K testing. The DNC runs for the Core i3-2120 in *Assassin's Creed: Origins* were actually to do with how the benchmark performed. It genuinely did complete each run, but there was massive lag in loading textures, with wire meshes and HD textures taking an eternity to load—if at all. As the scenes became more complex, frame stuttering occurred, and although the frame rate remained high, it was clear that this isn't how the benchmark was meant to run.

INTEL CORE I7-8700K CLOCK SPEEDS & CORE SETUP

	Cinebench R15 Single-Core (Index)	Frequency/Active Cores
Intel Core i3-2120	80	3.1GHz/2+2
Intel Core i5-3570	132	3.2GHz/4+0
Intel Core i7-4790K	169	3.7GHz/4+4
Intel Core i3-8100	154	3.6GHz/4+0
Intel Core i7-8700K	190	N/A

SSD RAID 0 TESTING (2X SAMSUNG 960 PRO 512GB)

	CDM Seq Read	CDM Seq Write	CDM Random 4K Read	CDM Random 4K Write
Intel Core i7-8700K	3,456 MB/s	2,170 MB/s	49 MB/s	138.3 MB/s
AMD Ryzen 7 2700X	3,066 MB/s	2,006 MB/s	49 MB/s	123.7 MB/s
Intel Core i9-7900X	3,157 MB/s	2,165 MB/s	43 MB/s	183.32 MB/s
AMD Ryzen 7 Threadripper 1920X	5,912 MB/s	3,859 MB/s	38 MB/s	180.2 MB/s

NVIDIA GEFORCE GTX 1060 3GB (AVG FPS) @ 1080P

	Assassin's Creed Origins	Middle Earth: Shadow of War	3DMark Sky Diver
Intel Core i3-2100	58 (DNC)	74	17,091 (index)
Intel Core i5-3560K	54	74	21,563 (index)
Intel Core i7-4790K	58	73	26,943 (index)
Intel Core i3-8100	57	74	23,182 (index)
Intel Core i7-8700K	58	72	32,900 (index)

NVIDIA GEFORCE GTX 1060 3GB (AVG FPS) @ 4K

	Assassin's Creed Origins	Middle Earth: Shadow of War	3DMark Sky Diver
Intel Core i3-2100	24	26	11,310 (index)
Intel Core i5-3560K	24	26	15,632 (index)
Intel Core i7-4790K	24	26	17,098 (index)
Intel Core i3-8100	24	26	16,581 (index)
Intel Core i7-8700K	24	26	17,033 (index)

NVIDIA GEFORCE GTX 1080 TI 11GB (AVG FPS) @ 1080P

	Assassin's Creed Origins	Middle Earth: Shadow of War	3DMark Sky Diver
Intel Core i3-2120	65 (DNC)	92	19,286 (index)
Intel Core i5-3570	65	119	27,547 (index)
Intel Core i7-4790K	84	120	37,042 (index)
Intel Core i3-8100	67	122	29,496 (index)
Intel Core i7-8700K	98	123	49,871 (index)

NVIDIA GEFORCE GTX 1080 TI 11GB (AVG FPS) @ 4K

	Assassin's Creed Origins	Middle Earth: Shadow of War	3DMark Sky Diver
Intel Core i3-2120	54 (DNC)	51	11,066 (index)
Intel Core i5-3570	49	51	16,303 (index)
Intel Core i7-4790K	51	51	22,582 (index)
Intel Core i3-8100	52	51	17,340 (index)
Intel Core i7-8700K	52	51	31,383 (index)

Attempts to clone editor Alan have gone terribly wrong. We've confined these clones to an exact copy of the *Maximum PC* office. They seemed fooled, for now... *By Neil Mohr*

VIRTUAL

VIRTUALIZATION is almost as old as our beloved integrated silicon chips. At the beginning of the 1960s, there had been two major issues. Firstly, many individual mainframe models were bespoke, so incompatible. The other issue was that as integrated processors became more powerful, institutions wanted to implement flexible "timesharing" between multiple users.

IBM dismissed this multi-user fad—batch processing was definitely the future! But in 1963, it lost a large MIT contract to General Electric. Realizing its huge mistake, it developed the general-purpose S/360 architecture, which could be implemented on a wide range of compatible systems. In 1965, IBM released the S/360-76, the world's first mainframe to support virtualization. And the rest is very much history.

The reasons of 55 years ago remain the same today. Virtualization is an ideal way of running multiple operating systems on the same hardware, sharing resources, and maintaining security through separation—building walls is good, sometimes. It's just you don't need a room-sized bit of tech anymore; even the most entry level of processors can happily churn out a bit of multi-machine virtual loving.

We imagine that almost every *Maximum PC* reader is packing a lean, mean virtualization machine at the heart of their mighty, throbbing computing rig. Toss in the right open-source software, and—boom!—you are ready to fire up everything from a Windows 10 clone, handy Android device, or dedicated gaming server to a classic-compatible Windows 3.1 PC. So, rev up those cores, we're going virtual machine racing.

You've probably already heard of VirtualBox—it's the best known virtualization software out there, alongside VMware and perhaps Xen, KVM, and Hyper-V. The truth is, VirtualBox is just one of many virtualization methods, but being open source, easy to use, and over 11 years old, it's been propelled to the fore over paid-for proprietary options such as VMware or Microsoft's Hyper-V.

Once you've got your real Windows install just right, the last thing you want to do is mess things up installing dirty software on it. VirtualBox enables you to try out software, run random services, and even share your system with so-called "loved" ones, who just want to take advantage of your fancy rig.

VirtualBox is regularly updated, and you can grab the latest release from www.virtualbox.org. At the time of writing, v5.2.8 was out, loud and proud. So, we're going to take a deep dive into setting up, using, and configuring VirtualBox, and the neat tricks you can use to access and use your VirtualBoxes.

01: Oracle might not be the most loved company in the world, but we've got to love VirtualBox.

01 VirtualBox helps you build a virtual wall, keeping all invaders safely trapped, running inside a seemingly happy real system of their own. Once you've installed and fired up VirtualBox, let's get some lingo down. Your real PC running VirtualBox is called the host; VirtualBox is technically the

BEING

The main "Machine Tools" interface displays all your virtual machines.

hypervisor; and any virtualized systems it runs on the host PC are called guests, and are your virtual machines (VM).

If you haven't done so already, start VirtualBox. The main "Machine Tools" interface displays all your created virtual machines. Click the top-right "Global Tools" button to list all the created virtual media drives, which are the virtual drives you create to store your virtual machines' OSes and data. Technically, these are reusable and portable, plus you have snapshot and history capabilities, hence the separate manager.

Before we create anything, though, consider where your VMs will be stored—both the small VM configuration (.vbox) files, and the potentially large virtual media drives (.vdi), and their snapshots. If you're running a smaller SSD boot drive, you'll likely want to tell VirtualBox to store these on a data drive. A basic Windows 10 image can happily hit 20GB, and if you end up playing with a few VMs, each could easily be 4GB. Specify a new location by selecting "File → Preferences → General Default Machine Folder location." While you're here, you might want to adjust "Update settings" to check once a month rather than daily.

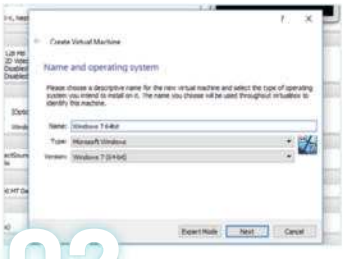

02

Click the "New" button top-left, or through the "Machine → New" menu. This launches a dialog that steps you through the creation of a new VM. We'll go through this, pointing out some of the finer points to keep an eye on. Off the bat, you need a unique name for your VM—you can't have two VMs called "Windows." It's also semi-intelligent when you select a "type," in that it selects some default settings, such as memory capacity, based on what it knows is needed.

02: Ensure you set your bits correctly; if you're installing a 64-bit OS, that's what you need to select.

03: Good grief! A genuine reason to be running 16GB on our PC!

04: With dynamic drives, there's no reason not to go wild and give your VMs plenty of drive space.

05: VirtualBox will play nice with your real processor's time, but you might want to limit the cap just in case.

06: It says "3D Acceleration," but don't expect VirtualBox to go playing any DirectX games, now.

The "32-bit/64-bit" setting here is important. A 64-bit host can run both 32-bit and 64-bit guests; 32-bit hosts (which aren't the best idea anyway) are limited to 32-bit guests. You also need to ensure the install media (we'll come to this shortly; usually a DVD ISO image) needs to match—so, a 64-bit ISO image for a 64-bit guest—otherwise the installer is likely to error at some point.

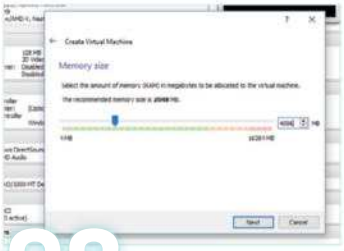

03

VirtualBox wants to know how much memory to assign to the VM. You can't allocate more than the host machine, but if you can, 4GB for a Windows Vista, 7, 8, or 10 guest is advisable. Next is the virtual hard drive. You can reuse an existing image with the bottom option, but we're assuming this is your first run, so select "Create a virtual disc now," then click "Create."

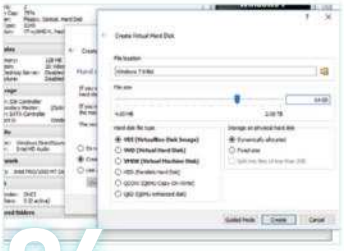

04

Click the "Expert" button, and you can edit all the options at once. The default settings should be adequate. The file type can be changed if you need to move it to another hypervisor. Fixed size drive images offer minor speed increases, but with obvious storage implications. With a Dynamic allocated drive, you can happily increase the upper file size without penalty, because it starts off sub-20MB in real size terms. Finally, you can click the yellow folder icon to check the storage location.

At this point, the VM is created, and you can start it, but we haven't allocated any boot media yet, so it won't get very far. Click "Settings" to discover the full capabilities—we'll go through the highlights in turn. Under "General → Advanced," switch "Shared Clipboard" and "Drag'n'Drop" to "Bidirectional." Once a guest operating system is up and running, you'll be able to install a pack called Guest Additions, part of which offers up the ability to drop files in and out of the guest, plus share the clipboard.

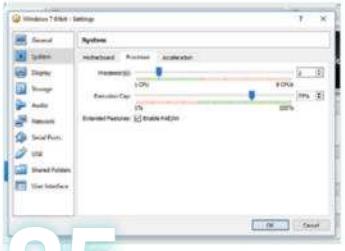

05

The "System" section controls vital motherboard-level and memory settings. "Chipset" should be set to "ICH9" when using Vista or more recent versions of Windows, because it supports PCI Express and modern interrupts—as a matter of interest, it supports up to 30 network adapters. Ensure "I/O APIC" is selected for virtual multi-core support. The "Processor" tab enables you to set the number of virtual cores and utilization cap. "PAE" is a processor feature for 32-bit processors to extend memory support beyond the 4GB limit. Some operating systems (Ubuntu, for example) require this to be enabled.

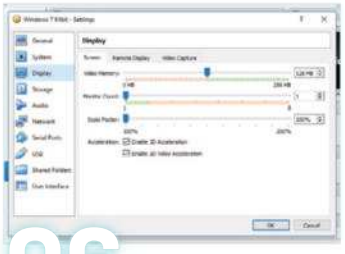

06

Display settings have advanced nicely in VirtualBox. Memory is allocated from the host's system memory; typically, 128MB is a reasonable amount. Multiple virtual displays are supported, while 2D and 3D acceleration are possible via Guest Additions, which we'll cover shortly. Remote Display is a handy feature enabling Remote Desktop access—we cover this in depth toward the end of the feature.

HARDWARE ACCELERATION

For x86 processors (other processors, such as ARM, have their own technologies), there is a number of different hardware features that have been added over time to help accelerate virtualization of guests. It's worth covering these, because they offer an insight into the work the virtual host is having to do.

Virtualization software largely has two options when it comes to running guest software faster: full virtualization or paravirtualization. The full option, as you might guess, requires the host to recreate every aspect of the virtualized hardware from the full instruction set, I/O operations/devices, interrupts, and memory. For x86 code running on an x86 machine, instructions can largely be run directly, but any instructions or I/O access that violate the virtual machine have to be trapped and translated—this slows everything down.

To help speed up virtualization, AMD and Intel introduced a range of processor features. The initial ones, called Intel-x and AMD-V, provide instructions to rapidly switch the guest to a privilege state to run trapped calls. AMD also added Rapid Virtualization Indexing, followed by Intel EPT, that virtualize memory page-tables lookups. Benchmarks show these increase kernel-level performance two to five fold, depending on the workload.

AMD-Vi and Intel VT-d provide an I/O memory management unit, also called PCI passthrough. It enables the guest OS to access real hardware seemingly directly, via a translation table on the host. Both processor and motherboard need to support the feature, and you may need to activate this in the BIOS.

Building on the I/O features, virtual interrupt support was added shortly after IOMMU; AMD, with its AVIC (Advanced Virtual Interrupt Controller), and Intel with its APICv. Both are designed to route virtual interrupts to the correct real CPU.

The latest development is on direct virtual graphics access. Intel has introduced GVT-d/g/s on the Iris Pro, and offers the potential for dedicated or shared VM access.

Making it happen first: the Athlon 64 with its AMD-V hardware support.

VIRTUALIZATION VS. EMULATION

Really, these are two sides of the same coin. Emulation is used to describe full-system virtualization, which also includes the guest's processor and all hardware. A VM created in VirtualBox will largely be using the same hardware as the host, so a lot of optimizations can be made that way. Usually with an emulator, the processor, graphics, and custom hardware controllers all have to be trapped and remapped to IBM-PC hardware.

Virtualization can also be done in different ways. You might have noticed (or possibly not) that Microsoft has been sucking up to Linux more and more in recent years. Not only is Microsoft using open source and Linux system internally (and admitting it)—it's offering Linux in its own store, so you can run it on Windows—but it's also bundling Linux distros for its Azure cloud service.

Part of the reason is that the Linux kernel has a virtualization system built in called KVM (no, not the monitor-sharing switch boxes). Bought by Linux powerhouse Red Hat years ago, KVM was mainlined into the Linux kernel back in 2007, and gives every release of Linux enterprise-level hypervisor capabilities at the lowest software level possible. Which is nice. It's supported by every architecture of note, such as ARM, PowerPC, IBM S/390, x86(-64), and Itanium IA-64. Just to highlight its power, Amazon announced it would be switching from the proprietary Xen system to KVM to power its monstrous Amazon Web Services business.

Of course, snobby Linux types can point back to the chroot jail system being in place as far back as 1979 (14 years before Linux was created), but it was in Linux use around 1999. Chroot is a basic system that "CHanges ROOT," root being the base OS folder. The chroot system enables Linux to run an OS-level application contained in its own software jail.

This system has recently spun out a series of innovations, such as containers and Docker, which build upon this base OS feature and have turned it into commercial applications for sharing, cloning, and reverting development and live environments.

07: Storage is important in VirtualBox. We won't go into snapshot management here, just the device settings. VirtualBox should use a SATA controller by default; there is IDE support for older OSes, plus SCSI, too. Booting for USB and NVMe devices requires EFI support under "System" to be enabled. It's not recommended that you enable the "Use Host I/O Cache" on a VM's drive controllers; VirtualBox has its own file system cache—a host system can end up allocating a lot of resources to caching a VM's read/write activities. It's also possible to limit individual drive bandwidth, but only via the command line, which is beyond this feature's scope.

Attaching your boot media can be done here. These days, if you fire up a VM without any boot media, VirtualBox asks you whether you want to select an image. Typically, this is an ISO file image—you select the optical drive icon, click the top-right optical disc image, and locate your ISO file. If you were doing this for a legacy OS using floppy drives, you'd do the same, but select the floppy image for the floppy controller.

We'll skip the "Audio" and "Serial" settings. Audio is largely there to be disabled if required, and if you need a serial port, you really should know what you're doing.

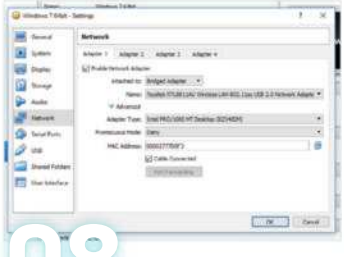

08: You can skip past "Networking," too (and it'll work perfectly)—the capabilities here are powerful, but way beyond what the average home user needs. At the most, if you want to run a network service, select "Bridged Adapter" in the "Attached to" pull-down. This ensures your VM gets its own IP address from your local router.

VirtualBox has full USB support, but it needs to be configured for it to identify host-inserted or even remote-inserted (when using Remote Desktop Protocol) USB devices. You have two options under the USB section: Use the green plus "+" icon to add a currently connected physical device to the filter list, or use the blue dot icon to create a general filter rule, including support for remote accessed devices.

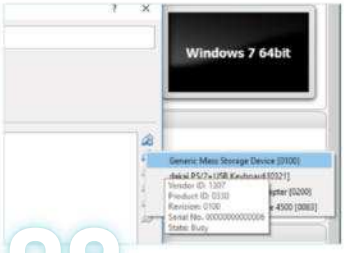

09: Finally, our favorite overlooked VirtualBox feature is "Shared Folders." When we've had issues getting drag and drop to work, the easiest way to get files on and off a VM (other than logging into your Google account) is to use this feature. You simply specify a folder on your host system, and this becomes exposed on the guest as a folder, too.

At this point, you should be primed and ready to launch a perfect virtual machine (VM) guest. You just need an OS install ISO file—or do you? It turns out you don't even have to do any of the work of installing operating systems these days—pre-built VM images are available, just waiting for you to download them.

In the past, we would have had to put in disclaimers about owning Windows licenses, but you don't have to worry for testing—just head to https://developer.microsoft.com and download a developer image of Windows. Fancy a more challenging operating system? Then Linux, in its many, many forms, can be snatched from www.osboxes.org/virtualbox-images. Of course, you can still install from scratch, but you don't really need to do that either, because www.turnkeylinux.org offers VM and ISO images of every important online service ready to run, from WordPress to OpenVPN.

Depending on the file, you need to do different things. An OVA file is a portable VM image; you import the "appliance" into VirtualBox via the "File → Import appliance" menu. A VDI file is a hard drive image—create your VM guest setting, and point the hard drive controller at this. Finally, ISO files are mounted as optical discs, and boot the VM ready to install and run.

07: A wide range of drive controllers are supported, including NVMe for plenty of speed.

08: There are overly comprehensive network settings; you'll likely only ever want to switch to Bridge mode.

09: Use the USB filter to help attach real USB devices to your VM for access.

As the VM guest first powers up, there is indeed a "BIOS"—hit F12 to enter this; really, it only offers boot device options. Depending on the OS and mode, if the VM captures the pointing device, you press right-Ctrl to swap back out. For Windows, MacOS, and Linux operating system guests, VirtualBox offers to install Guest Additions. Once up and running, you might need to select the "Devices → Insert Guest Additions CD" menu, and run this from within the VM guest. This installs extensions to the guest OS to enable extended features, such as shared clipboards, seamless mouse support, and on-the-fly resizing.

VirtualBox includes advanced snapshot controls over its drive images. Open the "Machine Tools" menu, and choose "Snapshots." This enables you to manage your VM images. "Clone" should be clear enough—it creates a duplicate image of an existing VM. Snapshots, however, create a delta image of an existing state, so you're able to revert back to the original snapshot at any point. Opening "File → Virtual Media Manager" enables you to see these in more detail, alongside being able to physically move them around your system, if you need to.

A very useful aspect of VirtualBox is its Remote Display Server—it enables multiple users to remotely log in via a Remote Desktop session, and control the VM—alongside its headless or detachable VM modes. This is a server that runs independently of the guest; you need to know the IP address of the PC VirtualBox is running on. To find this, in "Search" type "cmd" to open a command prompt. Type "ipconfig" and it's the listed IPv4 address.

Pre-built VM images are available, just waiting for you to download them.

A HEADLESS GUEST

To set up remote display, in the VirtualBox Display settings for the VM, click the "Remote Display" tab, and select "Enable Server." It's recommended that you change the default port—the port is how TCP-based programs connect to specific network services. Oracle recommends between 5000 and 5050, so let's say 5000.

Type "RDP" in "Search" on another Windows PC on the same network, and open Remote Desktop Connection. In the "Computer" box, type "<ipv4 address>:5000" and click "Connect." If you're on the same PC and want to test, use "127.0.0.1:5000"—the local loopback address. From a command prompt, you can type "mstsc -v <ipv4 address>:5000" and get the same result. Headless mode can be launched from the GUI by clicking the down arrow next to "Start." It launches the VM without a window, just for remote access.

A headless guest haunting our system, but being accessed via a remote desktop session.

DEDICATED SERVER

A fantastic use for a virtual machine is running a development system. Snapshots enable you to play with a guest VM, and rewind all your changes. As an example, we'll explain how to run a Turnkey dedicated WordPress server. This enables you to run your own full-fledged local WordPress server, and do whatever you like to it.

Head to www.turnkeylinux.org/wordpress and grab the VM .ova file—it's a reasonably light 304MB. Use the VirtualBox "File → Import Appliance" menu. Peruse the default settings—they should be fine—and click "Import." You might want to click "Settings → Network," and ensure Bridge mode is selected—it should be. Fire up the VM; it asks you to set root, MySQL, and WordPress passwords. As this is just for fun, "Passw0rd!" works well. You may want to skip the next options, as we're not fussed about security updates, since we're just running this locally as an example.

Keep an eye on the VirtualBox display, as it will show the local IP address when the WordPress server is up and running. Tap this IP address into your browser's address bar, and the template Turnkey WordPress website appears. There's instructions on how to log into the WordPress admin, and you set the password earlier, remember? It's also worth trying the Webmin interface you'll need to use: https://<ip address>:12321. Use root as the admin name, and again the password you set earlier.

Setting up a host of local dedicated servers is made a snap with VirtualBox.

MALWARE MACHINE

A good use of a virtual machine is to test suspect software. A complete setup would include a virtualized network, with any number of victim guest VMs, alongside a virtual gateway VM, enabling it all to be isolated from the real Internet. This would enable you to infect a guest, and get a full analysis of who/what it tries to contact on the real Internet, along with any additional network and local PC activity.

This type of work is beyond this feature, but you could use a VM to download suspect files, run them, and monitor what they get up to locally. Simply wipe the guest VM once you're done via the snapshot system. Just be aware that certain advanced malware can attack local systems through multiple networked vectors, so isolate the VM during testing by disconnecting its network capabilities.

Be aware that as malware has become more advanced, VMs have become vulnerable to attack. A small number of vulnerabilities have been discovered with VirtualBox that enable an attacker within a guest VM to break out of the hypervisor and attack the host. Such exploits are rare but do happen, so keep your system and VirtualBox up to date. Restrict network access with a properly configured firewall, run VirtualBox as a regular user, and monitor network and system activity.

SUBSCRIBE TO MAXIMUM PC

THE GO-TO RESOURCE FOR PC ENTHUSIASTS...

CHOOSE YOUR PERFECT PACKAGE

1

THE DIGITAL EDITION

ONLY $9.00

EVERY YEAR

Instant digital access on your iPad, iPhone, and Android device.

2

THE PRINT EDITION

ONLY $24.00

EVERY YEAR

Every issue delivered to your door for a fraction of the store price.

3
COMPLETE PRINT & DIGITAL BUNDLE
ONLY $28.00
EVERY YEAR

Get the print edition of *Maximum PC* delivered direct to your door, and instant access on your iPad, iPhone, or Android device.

On iOS & Android!

THE EASY WAYS TO SUBSCRIBE...

http://bit.ly/sub-maxpc

Or call 844-779-2822 (toll-free)

TERMS AND CONDITIONS Prices and savings quoted are compared to buying full-priced US print and digital issues. You will receive 13 issues in a year. You can write to us or call us to cancel your subscription within 14 days of purchase. Your subscription is for the minimum term specified, and will expire at the end of the current term. Payment is non-refundable after the 14-day cancelation period, unless exceptional circumstances apply. Your statutory rights are not affected. Prices correct at time of print and subject to change. Call center opening hours are Monday to Friday 8am–7pm, and Saturday 10am–2pm EDT. For full terms and conditions, please visit http://bit.ly/magtandc. **Offer ends: June 30, 2018.**

TECH PORN

1 PLANAR

There's a reason these headphones sound so good, and it comes from the planar-magnetic drivers nestled in each earcup. Capable of driving an immensely wide frequency response, the longer you wear them in, the better and warmer they sound.

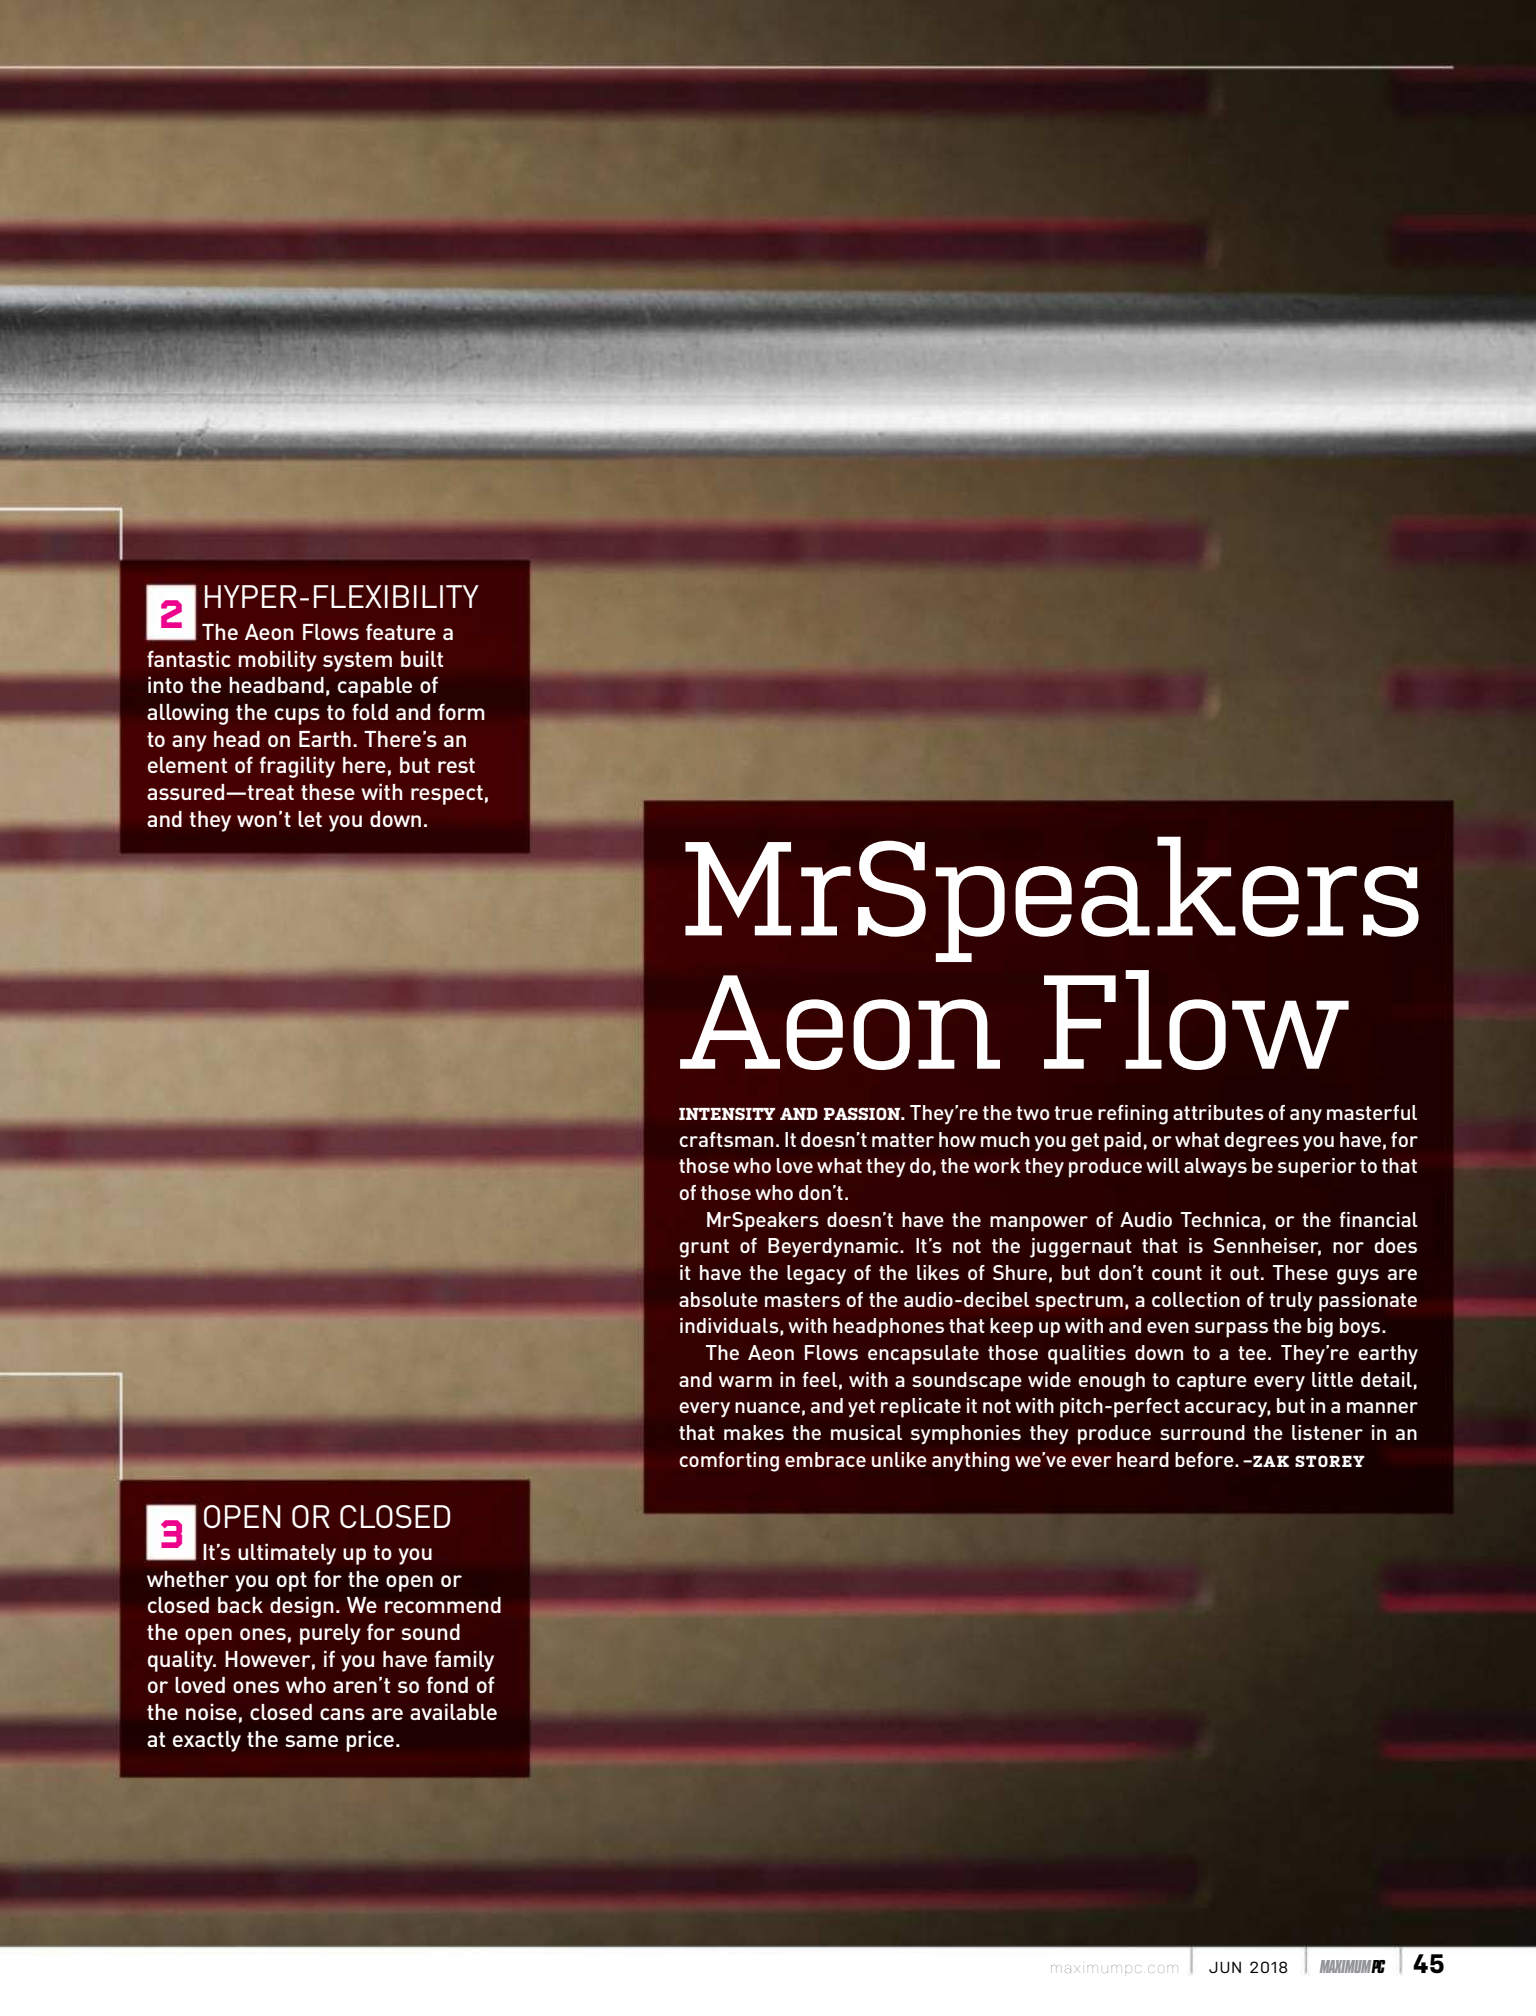

2 HYPER-FLEXIBILITY

The Aeon Flows feature a fantastic mobility system built into the headband, capable of allowing the cups to fold and form to any head on Earth. There's an element of fragility here, but rest assured—treat these with respect, and they won't let you down.

MrSpeakers Aeon Flow

INTENSITY AND PASSION. They're the two true refining attributes of any masterful craftsman. It doesn't matter how much you get paid, or what degrees you have, for those who love what they do, the work they produce will always be superior to that of those who don't.

MrSpeakers doesn't have the manpower of Audio Technica, or the financial grunt of Beyerdynamic. It's not the juggernaut that is Sennheiser, nor does it have the legacy of the likes of Shure, but don't count it out. These guys are absolute masters of the audio-decibel spectrum, a collection of truly passionate individuals, with headphones that keep up with and even surpass the big boys.

The Aeon Flows encapsulate those qualities down to a tee. They're earthy and warm in feel, with a soundscape wide enough to capture every little detail, every nuance, and yet replicate it not with pitch-perfect accuracy, but in a manner that makes the musical symphonies they produce surround the listener in an comforting embrace unlike anything we've ever heard before. **–ZAK STOREY**

3 OPEN OR CLOSED

It's ultimately up to you whether you opt for the open or closed back design. We recommend the open ones, purely for sound quality. However, if you have family or loved ones who aren't so fond of the noise, closed cans are available at exactly the same price.

REAL-TIME RAY TRACING

The next big leap in graphics processing is upon us—well, almost...

By Alex Cox

The Reflections demo looks stunningly realistic, as much as something featuring stormtroopers can.

WHAT'S OLD IS NEW AGAIN. Ray tracing, that stalwart tech that made the Commodore Amiga look graphically impressive (as long as you left it on for two days, and didn't want to render more than a sphere and a checkerboard) is making a comeback. The light-following rendering technology is on the way to making the move from multiple days of processing time to 60fps, and this means a massive change for previsualization in movies, the graphical fidelity of games, and possibly the entire future of GPUs.

Big names in rendering have jumped on board, such as Epic and its Unreal Engine, Unity 3D, and EA's Frostbite. Microsoft has pledged its support in a big way, with an entirely new set of rendering instructions integrated into DirectX 12, and Nvidia is working to integrate real-time ray tracing capabilities into its hardware with its own additional processing layers.

Real-time ray tracing is possible now, but when will it land on your desktop? What tech does it require, and how does it work? What impact will it have on tomorrow's games? And what will it cost? The answers are slightly more fuzzy than the shiny, reflective surfaces of your traditional ray-traced scene. We know partial answers—we know a lot about the ins and outs of how the tech works, and the ways in which Nvidia et al are cheating their way to reflective, refractive scenes without the noise and rendering overhead that ray tracing normally throws up. We know the tech that's powering the wow-worthy demos. But the end result, and ray tracing's future as a core part of PC gaming, isn't so clear. So buff up that GPU, polish your monitor, and prepare for a future where everything's slightly more reflective and much more true to life...

The post-apocalyptic world of *Metro Exodus* ain't shiny, but the shadows are quite incredible.

Pica Pica shows off global illumination and soft shadows.

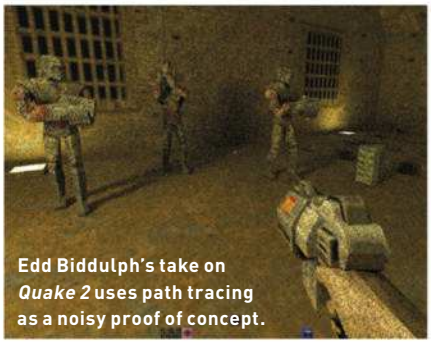

Edd Biddulph's take on *Quake 2* uses path tracing as a noisy proof of concept.

Remedy's demo shows off some impressive light rays and shadows.

WHAT IS RAY TRACING?

The clue is in the name. It's is a rendering process that traces rays. In the most basic job of the process, a ray is traced from the viewport or camera through each individual pixel of the screen, projecting out into the rendered scene. When it intersects with an object, it reflects back the precise color that individual pixel should be. Simple so far, but that color is determined by three factors: the amount and color of the light falling on the object, the reflectivity or material of that surface, and the position of the camera.

These, inevitably, break down into even further detail. To determine the amount of light falling on to a single pixel, the ray-tracing formula needs to know how far away the light is, how bright it is, and the angle of the reflecting surface relative to the angle of the light source, before calculating how hot the reflected ray should be. The process is then repeated for any other light source, including indirect illumination from light bouncing off other objects in the scene. Calculations must be applied to the materials, determined by their level of diffuse or specular reflectivity—or both. Transparent or semi-transparent surfaces, such as glass or water, refract rays, adding further rendering headaches, and everything necessarily has an artificial reflection limit, because without one, rays could be traced to infinity.

Even considering all the individual factors in play, every pixel is created by a simple linear equation. Simple math which, as you might imagine, explodes into a whole lot of not-so-simple math very quickly. A single frame might be rendered four times—once for specular reflections, once for diffuse, once for environmental reflections, and once for color—before being combined together. This is incredibly, unfathomably computationally expensive, which is why, for many years, ray tracing has generally been confined to slow-moving render farms; Pixar's *Monsters University,* admittedly running on 2013 hardware, took 29 CPU hours to render a single frame. Real-time graphics have instead had to fake it till they made it. From the early days of 3D, developers have used raycasting—a vastly simplified form of ray tracing, which only takes the initial light reflection into account—and rasterized it to the screen, then piled on effects, shaders, and efficient cheats to make the end result look realistic.

Ray tracing has a long history, but its sluggish speed has always made it merely a real-time dream. It's something programers have tinkered with, on and off—in mid-2017, coder Edd Biddulph uploaded a video showing a version of *Quake 2* he'd tweaked to use a path-tracing algorithm (essentially the same idea as ray tracing) that looked fuzzy and couldn't even manage a consistent 60fps on a Titan XP.

DXR, WTF

We're now nearing the top of the hill, and soon we'll tumble down the other side in a bundle of reflective, beautifully-lit realism, thanks to the combination of Nvidia's new

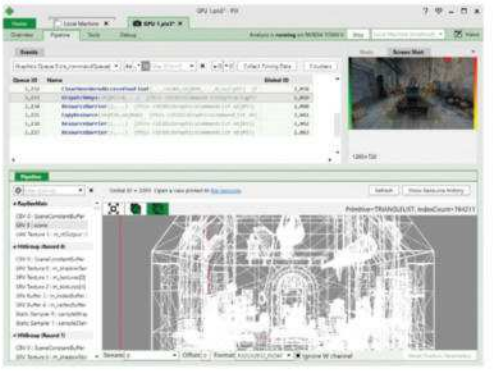

Microsoft's PIX performance tuning tool already includes support for DXR ray tracing.

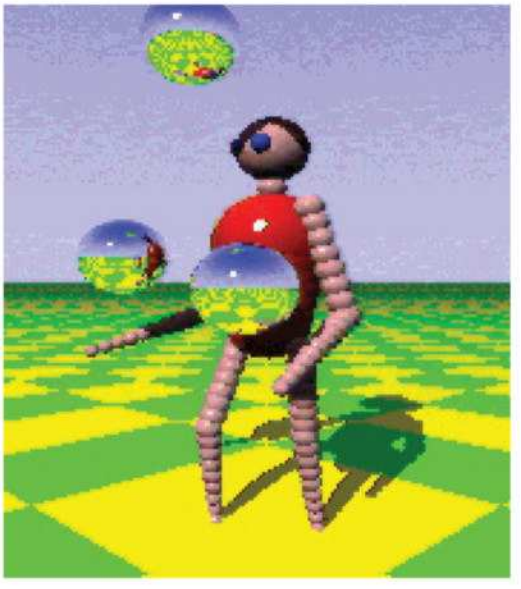

Ray tracing's consumer roots go way back to the Commodore Amiga's Sculpt 3D and beyond.

MAXIMUM VOLTAGE

Nvidia's next-gen Volta line of 12nm GPUs are incredibly rare at this point, existing only in the mid-2017 business-class Tesla V100 and the $3,000 Titan V, released to consumers last December. There's your first barrier to ray-tracing entry: Volta is seriously expensive, and if the crypto market has anything to say about it, even process-refined cards will command a high price.

Its architecture has a number of key advancements, beyond the compressed layout of the GPU core, that make it highly suited to real-time ray tracing. First, the tensor cores, which are specifically suited to processing, well, tensors—that is, functions of the co-ordinates of a space. You may have heard of tensors previously, thanks to Google's Tensorflow machine learning algorithm, and in the case of ray tracing, they're used to pony up AI to power the Optix denoising algorithm. This takes advantage of the unique abilities of the cores, which deal with mixed-precision FP16/FP32 calculations, to intelligently infer the position of misplaced rays in a scene in multiple dimensions.

Improvements in Volta don't end there. A massive amount of high-bandwidth memory, in the form of at least 12GB of GDDR5X, with a memory bandwidth of 652.8GB/s, makes throwing around the relevant ray-tracing data structures, and making them instantly accessible to the routines that need them, a trivial task. The natural increase in cores and transistors that comes from a jump in manufacturing process—the Titan V manages 21.1 billion transistors, and six graphics processing clusters with 5,120 CUDA cores—means a huge increase in performance; a single Titan V can provide 110 teraflops of power, against the Titan XP's 12 teraflops. All this, while consuming half the amount of power.

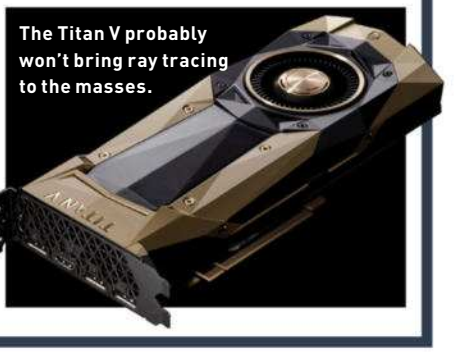

The Titan V probably won't bring ray tracing to the masses.

RTX tech, Microsoft's new support for ray tracing in DirectX 12, and the addition of raytracing support to popular rendering engines. Before we explore what we've seen, and the reasons you shouldn't get too excited just yet, let's look at the way DirectX Raytracing (which seems to be commonly referred to as DXR) processes ray tracing.

DXR isn't completely rewriting the way images are rendered. It's cleverly using a number of resources Direct X already works with, and giving ray-tracing engines the ability to hook on to them easily. Textures, buffers, and constants are all shared between the ray-tracing pipeline and other rendering dispatch types, with no conversion required between them, while shaders and general GPU synchronization tasks are open and application-controlled, allowing an efficient balance between rendering beauty and performance.

DXR's ray-tracing pipeline uses the same command lists and queues as the rest of its rendering tasks, meaning adapting renderers for ray tracing won't require a complete rewrite, and it fires off its jobs in grids, making the most of the parallel processing abilities of GPUs. There's still a little cheating, but this is mostly logical; DXR makes new objects available to the rendering pipeline, such as acceleration structures, which aid in quickly finding intersections, and shader tables, which quickly define the relationship between surfaces and light—reducing the amount of direct computation required per frame.

While the potential volume of computation has theoretically reduced, ray tracing—even under DXR—relies on the immediate availability of information. In this case, DXR implements what it calls state objects, which compile together every single shader of every single object that could possibly be hit by a ray, and defines the intersections and relationships between them, all ready to execute in an instant. Different scenes may have multiple state objects, and this puts a heavy load on graphics RAM, and asks a lot of a GPU's ability to process what needs to be seen in a scene. And Nvidia's RTX library, running between DXR and its Volta architecture, looks set to provide it.

RTX, OMG

One of the biggest problems in real-time ray tracing is noise, sometimes known as surface acne, as seen in the 2017 *Quake 2* demo. Flinging rays around from the viewport quickly, and within the limitations of floating point calculations, causes problems with the interaction of shadows and light rays, essentially causing them to intersect above or below the surface they are supposed to be reflecting from. Throwing enough rays to accurately cover the scene is a long way from being possible in real time. To fight this, we know that RTX includes an AI-infused denoising module, known as Optix, which helps to reduce jarring bald spots and graininess; doing this independently, rather than making it part of the initial optimization, is a much more efficient process than having to send out additional rays or laboriously recalculate intersections in order to fill the gaps. With

WHEN, WHEN, WHEN?

There's a good chance ray tracing will never become a primary tech for game developers—or at least, it won't be taking over for another decade or two. Doing it in real time demands a vast amount of processing power; consumer cards that can support even this early level of ray tracing will take a long time to hit our machines. In the meantime, rasterization is only going to get better. When the hardware to render everything with ray tracing does land, it probably won't look that different from the tech it's meant to emulate, making the overhead that ray tracing adds a pointless computational expense. Sticking with more efficient rasterization leaves room in the GPU for ever more complex physics and material simulation; making things a micron more shiny won't be worth losing the other tools a GPU can offer.

It's more likely that DXR, and ray tracing in general, will become a sideline tech for particular materials in particular situations—used sparingly, we could see ray-traced elements in games sooner rather than later. The majority of high-end 3D game engines have added support—but, at least on the Nvidia side, it's only viable as long as you use Volta or better, as Pascal doesn't have the tensor tech.

If you're in the movie business, expect to see DGX rigs dropping all around you soon. The value in being able to immediately render and see the results of visual effects, rather than waiting for servers to return rough scenes, will make CG movie-making more efficient and far more beautiful.

Nvidia's DGX platform could do for VFX what it's done for AI.

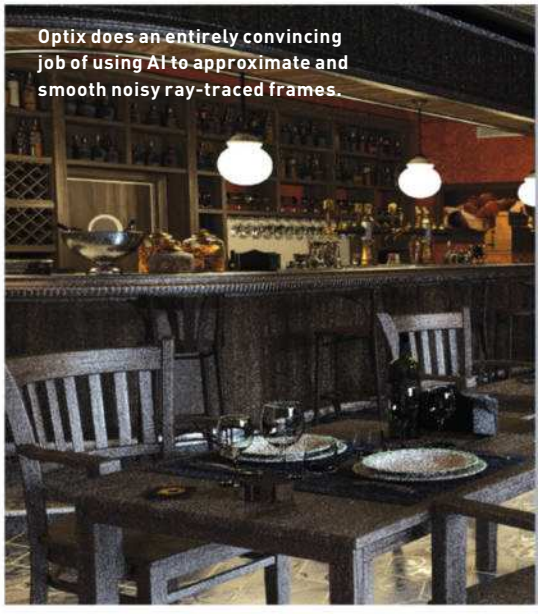

Optix does an entirely convincing job of using AI to approximate and smooth noisy ray-traced frames.

Nvidia's RTX tech demos highlight shadows, ambient occlusion, and reflections.

denoising, ray-tracing routines can send as few as two rays per pixel, and still achieve a solid result.

RTX is also where the critical management processes for shaders, state objects, and acceleration structures live, and the fact that it's inextricably linked with Volta means there's certainly a little more AI trickery and three-dimensional calculation going on, as well as an extreme amount of cache abuse. Accelerating ray tracing comes from doing as little calculation as possible—intelligently caching the intersections of one frame then doing simple transforms to predict where they'll land on the next, while not technically precise, is a great way to pull off realistic-enough ray tracing without breaking the GPU bank. By way of demonstration, tech demos have included automatic lightmap baking, presumably to fill those DXR shader tables ultra efficiently.

Much of what we know about RTX is based on what we've seen, or an educated prediction. Nvidia is, understandably at this early stage, keeping quiet about the exact processes within the RTX layer, although we know plenty about Volta (see "Maximum Voltage" pg. 49 for more), and Microsoft has been plenty open about how engine developers can interact with DXR. Given that higher-level support, the specifics don't matter too much at this stage, as long as they work, and they do—to an extent.

THE SHOWCASE

GDC 2018 marked a big turning point for the idea (and the execution) of real-time ray tracing, as every party even slightly involved with the tech started flexing their shiny oiled-up muscles so that you could really see their specularity. Nvidia's key demo, pulled off in Unreal Engine 4 with a fair amount of help from the geniuses at Industrial Light and Magic, was a dynamic *Star Wars* scene—aptly titled "Reflections"—starring a couple of stormtroopers and the chrome-armored Captain Phasma. To watch, knowing that what you're seeing is ray tracing calculated on the fly, it's incredibly impressive, particularly the AR-controlled live exploration of the scene pulled off on stage, but it's important to examine exactly what was included in terms of materials: It's all synthetic.

Everything in the 24fps Reflections demo is appropriately shiny, showing off some mind-blowing specular reflections, but the plastic and chrome combination, in what is a geometrically straightforward environment—three boxes, essentially, with a little set dressing, and a couple of ray-tracing favorite spheres for good measure—goes some way to pulling off an impressive demo without overtaxing the tech as it stands. To quote Epic Games' art director Jerome Platteaux, in an interview with Starwars.com: "We designed materials with proper shading response, so they behave naturally to light." It turns out, also, that the walls of the demo were rasterized, reducing the processing load considerably.

Similarly slick was EA SEED's demo of Pica Pica—a, well, thing with AI-controlled robots. We're not entirely clear

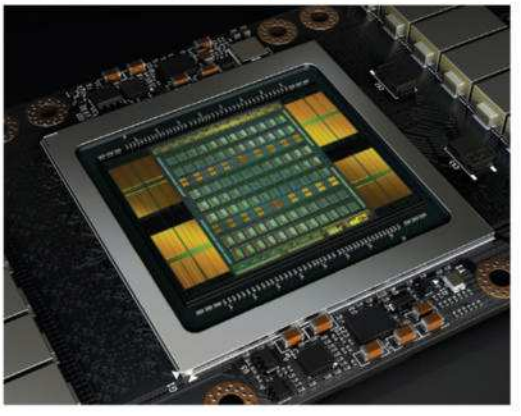

The tensor cores of the V100 chip enable the AI denoising algorithms of RTX.

RENDER IT YOURSELF

Although the tools to get started with real-time ray tracing are pretty much there right now, with DXR support added to Unity 3D, Unreal Engine, and the like, you're not going to be able to take advantage of them yet. But if you're interested in seeing how real-time ray tracing works on current hardware, and want an example of how valuable denoising is going to be, you can try some old-school tracing now.

Otoy's Brigade engine (http://otoy.com) is a real-time path tracer that purports to run on current GPUs—and it does, albeit in a very noisy way. You can try an online demo, running through a version of Otoy's OctaneRender package, at http://render.otoy.com/includes/xio_launch.php to see what the fuss is about and why denoising matters. There are multiple scenes that you can pan, zoom, and twist around; as you move the scene renders (slowly), starting with a limited number of paths traced, and increasing over time.

If you want to go really old-school with ray tracing, there are plenty of packages that can pull it off given a little (or a lot of) rendering time. These range from the nearly 30-year-old POV-Ray (www.povray.org) up to Pixar's pro rendering engine, Renderman (http://renderman.pixar.com), which is free to try for non-commercial use. The latter offers lots of documentation, so it's a great way to get started, and the former has a massive pedigree, and while it's long in the tooth, it's based on classic rendering principles, which stand up to this day.

on what Pica Pica is, but the demo, from EA's experimental wing, showed off neat refractions, new techniques for ambient occlusion and global illumination, and—wouldn't you know it?—a couple of chrome spheres, and some seriously specular reflections. The promise of cute little AI-controlled dudes certainly suggests that Pica Pica is a title specifically developed to take advantage of both the graphical muscle of next-gen GPUs and their potential to host neural networks. Again, though, it was a geometrically simple scene.

There were less shiny demos on offer at GDC, too. Remedy Entertainment pushed out an unnamed demo, essentially a fly-through of a scene, which upped the geometry and mixed a number of different materials together—but, we presume, dropped the ray-tracing resolution, because the whole thing looked a lot less polished, and not just in the diffuse versus specular sense. 4A Games offered up a look at *Metro Exodus* running with selective ray tracing implemented, a heavily HDR scene with a lot of complex geometry, though you'd have to squint hard to see the difference.

THE LIMITATIONS

Though companies have been keen to show off what they've achieved, it's obvious that the tech remains limited at this stage. The *Star Wars* demo, even with all its optimizations and fudges, was run on an Nvidia DGX Station, packed with four Tesla V100 cards rocking a total of 64GB of VRAM. To buy one, you'd need to empty your wallet to the tune of $69,000—and Nvidia's higher end graphics supercomputers are even more expensive, with the DGX-1, rocking eight Tesla cards and a mere 5,120 CUDA cores, costing $150,000.

So this won't be landing in your home yet. At least not in its current form. But there's hope on the horizon, with various companies looking at ways to fudge ray tracing on to less bank-breaking hardware. AMD is working on spinning RadeonRays and ProRender—which combine CPU and GPU processing—into a real-time ray-tracing and rasterization mixed tech that suits its 3D professional partners. Imagination Technologies has even been working on its own Vulkan-compatible real-time ray-tracing tech on its mobile PowerVR GR6500 GPU. Intel has been poking ray tracing with a stick for a number of years. And the denoising tech of RTX, the proof that it works and can turn grainy filth into beautiful surfaces, should inspire a lot of lateral thinking in the coding community. Even DXR throws us a tiny bone, including shader-based fallback for faking ray tracing without the need for RTX (or similar) tech. Things are gonna get more shiny.

We're not sure where the future of Nvidia's consumer cards lies. If it sees ray tracing as an essential—and the company seems to be cuddling up with Microsoft on DXR—future Volta releases may feature a revised architecture that prioritizes denoising over the general-purpose (and expensive) tensor cores in current cards. But that's a theory, and real-time ray tracing is some way from mainstream just yet.

XBOX
THE OFFICIAL MAGAZINE

ISSUE 213 MAY 2018
WWW.GAMESRADAR.COM/OXM
XBOX
THE OFFICIAL MAGAZINE
MOST WANTED!
ROCKSTAR GAMES PRESENTS
RED DEAD REDEMPTION II
UNMISSABLE INFO ON ROCKSTAR'S WILD WESTERN
CO-OP CAPERS!
STRANGE BRIGADE
We go on an exotic safari into danger with Rebellion
HANDS-ON VERDICT!
FARCRY5
23 TOTALLY INSANE THINGS YOU NEED TO TRY
Alien arachnids, zombies and gonads? Ubisoft takes a turn into madness...
NEW INFO!
ANTHEM
How a delay could make this a mech-suited masterpiece
BATTLE ROYALE!
DARWIN PROJECT
Is this the first PUBG-killer on Xbox One?
EXCLUSIVE INTERVIEW!
BIOMUTANT
Mutations, abilities and the huge open world detailed

The essential magazine for all things Xbox

NEW ISSUE ON SALE NOW

ORDER DIRECT FROM MYFAVOURITEMAGAZINES.CO.UK/GAMING

HOW TO

STEP-BY-STEP GUIDES TO IMPROVING YOUR PC

TIP OF THE MONTH

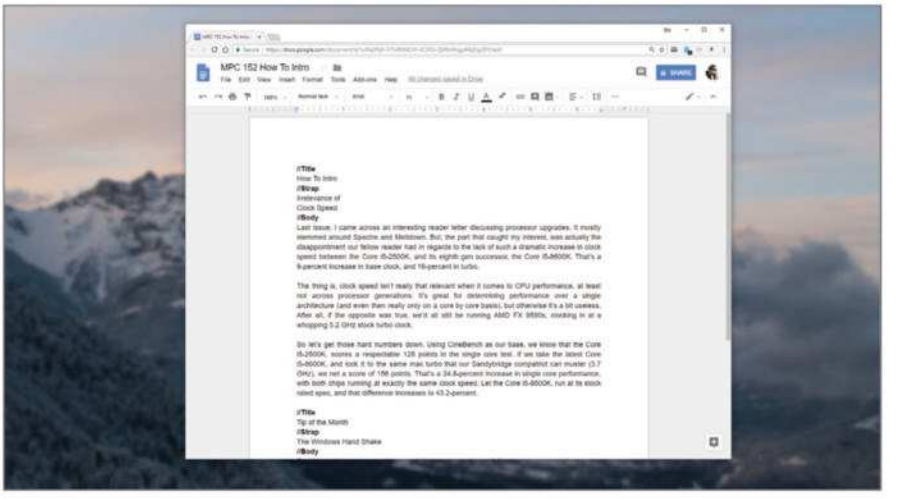

THE WINDOWS HANDSHAKE

Ever been driven insane by Windows 10's handshake feature? Shook the mouse too fast and had all your application panes minimize down to the Taskbar, only for you to have to painstakingly reopen them all again? Don't worry—there's an easy way out. Do exactly the same mouse shake again, and all those same application panels will open up to exactly the same position. $Free

MAKE – USE – CREATE

56 Protect against unexpected power loss on a Pi

62 Check your speed after Meltdown Spectre fixes

68 AMD's 12nm process makes its debut in our Build It section

ZAK STOREY
DEPUTY EDITOR

IRRELEVANCE OF CLOCK SPEED

Last issue, I came across an interesting reader letter discussing processor upgrades. It mostly stemmed around Spectre and Meltdown, but the part that caught my interest was the disappointment our reader had in regard to the lack of a dramatic increase in clock speed between the Core i5-2500K and its 8th Gen successor, the Core i5-8600K. That's a 9 percent increase in base clock, and 16 percent in turbo.

The thing is, clock speed isn't that relevant when it comes to CPU performance, at least not across processor generations. It's great for determining performance over a single architecture (and even then, really only on a core by core basis), but otherwise it's a bit useless. After all, if the opposite were true, we'd all still be running AMD FX 9590s, clocking in at a whopping 5.2GHz stock turbo clock.

So let's get those hard numbers down. Using Cinebench as our base, we know that the Core i5-2500K scores a respectable 125 points in the single-core test. If we take the latest Core i5-8600K, and lock it to the same max turbo that our Sandybridge compatriot can muster (3.7GHz), we net a score of 156. That's a 24.8 percent increase in single-core performance, with both chips running at exactly the same clock speed. Let the Core i5-8600K run at its stock rated spec, and that difference increases to 43.2 percent.

↘ submit your How To project idea to: comments@maximumpc.com

IFIXIT presents:

AUTOPSY

THIS MONTH WE DISSECT...

About iFixit

iFixit is a global community of tinkerers dedicated to helping people fix things through free online repair manuals and teardowns. iFixit believes that everyone has the right to maintain and repair their own products.
To learn more, visit www.ifixit.com.

All 32 IR sensors laid bare.

HTC Vive Pro

The Vive Pro features dual 1440p screens developed by Samsung.

BACKGROUND

With the lightning pace of tech releases, it's hard to believe nearly two years have passed since our first Vive teardown. Today, HTC brings you the first headset in what's bound to be VR's v2.0, and we're bringing you the pieces.

MAJOR TECH SPECS

- Two 1440p AMOLED displays for combined res of 2880x1600
- 90Hz refresh rate
- Dual front-facing cameras, dual mics, and headphones
- Accelerometer, proximity, gyroscope, and IPD sensors
- 360-degree headset tracking via up to four Lighthouse 2.0 base station IR emitters
- Wireless with the not-yet-released Vive Wireless Adapter

KEY FINDINGS

- We haven't even started disassembly, and already we're seeing some repair-friendly features. Spring contacts and standard Phillips screws connect the modular headphones to the headset, making for easy removal and replacement. HTC even tells you how to remove them. We peel off a couple foam pads and unsnap a plastic compartment cover to find the end of the Vive Pro's fancy custom cable. Although a completely wire-free experience is not yet available, HTC managed to fit all the power and data lines into a single cable that routes neatly up and over your left ear.
- Wading through the sea of divots in the front cover, we find a couple (or four) that hide Phillips screws. While we're not huge fans of hidden or tamper-evident screws, they're much better than glue or proprietary heads. After a bit of work with an opening pick, the outer shell splits down the middle and slides off. The Vive Pro's unsettling new gaze makes us feel bad about peeling its face off—sorry, Johnny Vive Five.
- Inside the covers, we spot familiar IR-friendly windows. All these purply dots are a bit like sunglasses for the IR sensors beneath—helping to filter out visible light, so the sensors can more easily spot the IR beams from the Lighthouse base stations. Much like on the original Vive, there's a spiderweb of photodiodes just under the sleek surface of the Pro—32 in total. Unlike the Oculus Rift, this headset is looking for IR light—not emitting it.
- Where the original Vive kept nearly all its onboard silicon on a single board, the sequel uses two smaller boards. One services the displays, I/O, and one or two other things. While we are excited about the chips, we're a little bummed that the ports are still soldered to a board. Non-modular, high-stress components like ports can be a pricey repair. This board also hosts the linear potentiometer that formerly lived on its own FPC. That small slider tracks where you position the displays when making IPD adjustments.
- Repairability Score: 8 out of 10 (10 is easiest to repair). Disassembly with standard tools is straightforward, non-destructive, and free of booby-traps. Newly added earphones are completely modular, and come with instructions for removal and installation. Standard Phillips and Torx screws secure most components. High-wear comfort pads are fastened with Velcro. Compatibility with existing controllers and base stations means you may already know how to repair the accessories. Adhesive is used sparingly, securing lenses, mic, and sensor arrays. This is an unusually complex device with a lot of delicate bits, and HTC does not provide a service manual—so use extra caution if attempting a repair.

Protect Against Power Loss on a Pi

YOU'LL NEED THIS

RASPBERRY PI
We use a Pi Zero in this tutorial.

UPS
Buy an Alchemy Power UPS from www.amazon.com.

YOUR RASPBERRY PI is particularly sensitive to power outages. Added to the inconvenience of suddenly losing access to your Pi or half an hour's work, there's the possibility your shutdown will corrupt your microSD card, forcing you to wipe it and start again from scratch.

This is where a battery rides to the rescue. It provides your Raspberry Pi with an independent source of power when the main supply is interrupted, whether through a power outage, tripped MCB, or simply by accidentally yanking out the power cable. While it is possible to jury rig your own battery supply, why go to all that trouble when a ready-made uninterruptible power supply (UPS) is waiting for you? –NICK PEERS

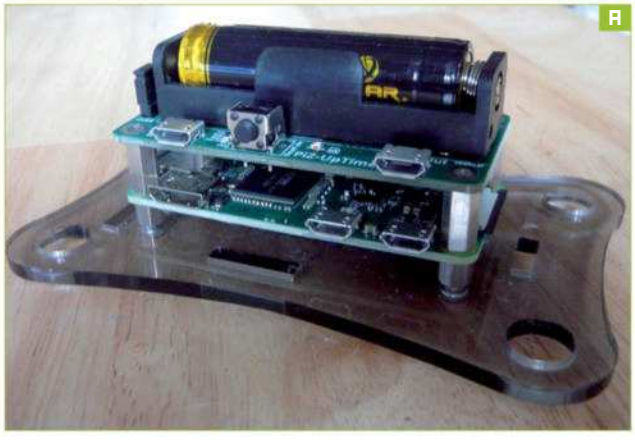
A

1 YOUR OWN PI UPS

The solution comes from the folk at Alchemy Power, who have developed two UPSes for Pi and Pi Zero users. Its main solution is the Pi-UpTimeUPS, which works with any Pi that has a 40-pin GPIO header, including the Zero. However, its form factor is aimed at the Pi, and it comes with a prohibitive price tag: $50 from Amazon.

» A less powerful option, but one that can still provide over an hour of emergency power, is the PiZ-UpTime. This costs $29 from Amazon, and is the unit we used in this tutorial. Although optimized for the Pi Zero, it works with the Pi 2 and Pi 3, too, albeit for a shorter period than the Pi-UpTimeUPS, which can deliver over 10 hours via its two batteries.

» Both models come as a single circuit board with a battery holder [Image A]. They can both be connected to your Pi via its GPIO pins (Pi Zero users need to solder or add these themselves), and both offer pass-through pins, so you can use them in conjunction with other add-ons, secure in the knowledge they've been designed to supply power to both your Pi and any add-ons in the event of power failure.

» Of course, it's not always practical to fit GPIO pins to your Pi Zero, which is where the PiZ-UpTime scores points. Not only can it be plugged into a GPIO pin header, but it can also be connected relay-style between your power supply and the Pi Zero using the supplied micro-USB cable.

2 PHYSICAL INSTALLATION

Power down your Pi, unplug the power supply, then fit the UPS board to your Pi, either via the GPIO pin header, or by unplugging the power from the Pi Zero, and connecting the PiZ-UpTime to the Pi Zero's power socket using the supplied cable.

» Once done, connect the battery—these tend to be supplied with some charge, so the moment you install it, your Pi leaps into life. Be sure to shut down safely from here — `ssh` into your Pi if necessary, and issue the following:

```
$ sudo shutdown -h now
```

» Now connect your Pi's power supply to the "In" port on the Pi-UpTimeUPS or PiZ-UpTime, so it can start charging the battery as well as provide emergency power when required. This follows the standard CCCV charging method, except when the battery has been depleted, in which case, low-current charging is first used to bring the battery up to a level where it can be safely CCCV charged.

» As you'd expect, the board has onboard circuitry to prevent both overcharging of the battery as well as deep

SOURCE THE RIGHT BATTERY

Batteries aren't supplied with either of the Alchemy Power UPSes we've recommended, and it's crucial that you get the correct ones. The Pi-UpTimeUPS uses 18650 non-protected batteries. There's room for two, but you can get by with just one. Recommended capacities are 2,000–3,500mAh per battery. In return, the UPS can deliver both 5V power via a micro-USB port (using pin 2), and 3.3V (700mA) of power via the Pi-UpTimeUPS itself.

In contrast, the PiZ-UpTime uses a single 14500 Li-ion battery, capacity 700–2,000mAh. Don't be fooled by their appearance—they look like regular AA batteries, but they're not, so don't be tempted to use standard batteries. Both are available from a range of sources, including eBay and Amazon. Remember, the higher the capacity, the longer your Pi (and any connected add-ons) can run on battery power—we got well over an hour out of an 800mAh Xtar-branded battery that we purchased on eBay ($6).

It's worth noting that the unit will perform a hard shutdown when the battery's capacity runs dangerously low, thus prioritizing the battery's health over a few extra seconds of power.

UNDERSTAND THE LEDS

The PiZ-UpTime and Pi-UpTimeUPS come with multiple LEDs to help you fathom what's going on. There's an orange LED that indicates the battery level. This starts off practically invisible, but glows brighter as the battery level drops, to provide an early warning system. You'll also find a red LED that lights up after the hardware reset switch has been pushed, to indicate the power reboot sequence.

Three more LEDs can be found along the side. The blue LED lights up when running on battery to indicate low battery (under 3V of power). The green LED indicates mains power is connected—when it is, the blue LED switches function to indicate that the battery is being charged. When charging is complete, the yellow LED lights up instead.

It's possible to disable these LEDs using a jumper on the board, which can save you a bit more power. Of course, this might be practical if you've connected the board via the GPIO pins and have the shutdown scripts in place. If you're forced to connect the PiZ-UpTime to your Pi using the supplied micro-USB cable, you need to shut down the Pi Zero manually in the event of power failure, which might be an issue if, for example, you suffer an outage and have no network access. Thankfully, if you choose a reasonable battery, you'll have plenty of time to get things running again before you need to panic.

discharging—a hard shutdown is performed before the latter can occur. There's also a temperature monitor that can shut down the charging process if it becomes too hot. Note the reset button on the front or side of the board. Use this after a controlled shutdown to restart your Pi without first having to unplug the power or disconnect the battery.

» The battery level is monitored on GPIO pin 26—a jumper on the board can be set to disable this if the pin is in use elsewhere. However, if you want to be able to perform a controlled shutdown when the battery is low, you need this connected. This is the biggest drawback of the PiZ-UpTime's USB connection option, as it can't provide this functionality unless connected via the GPIO pins. Instead, use the LEDs to keep an eye on things (see box above).

» To make use of the shutdown scripts [**Image B**], first update the underlying OS powering your Pi, typically like so:

```
$ sudo apt-get update
$ sudo apt-get upgrade && sudo apt-get autoclean
```

Next, create a folder and download the required files:

```
$ mkdir ups
$ cd ups
$ wget http://alchemy-power.com/wp-content/uploads/2017/06/GPIO-shutdown-sample.zip
$ wget http://alchemy-power.com/wp-content/uploads/2017/08/GPIO-crontab.zip
```

» Next, unzip the two Python scripts:

```
$ unzip GPIO-shutdown-sample.zip && unzip GPIO-crontab.zip
```

» When the first script is run, it sleeps for 10 seconds before checking the battery levels again. If the battery levels are low, it triggers a controlled shutdown within two minutes. You can see how the script works by examining it in Nano, where an explanation accompanies each command:

```
$ nano GPIO-shutdown-sample.py
```

» If you want to lengthen the battery level check, you need to change the "schlaf" value as shown in the script. The script also displays a time stamp on screen every time it wakes to check the battery level. You can disable this by commenting out line 78 as follows:

```
# print("Time is %s " % (time.ctime()))
```

» If you want to skip the two-minute warning when battery levels run low, change the "subprocess.call" line:

```
subprocess.call("shutdown -h now &", shell=True)
```

» Save any changes and exit. Now test the script:

```
$ sudo python GPIO-shutdown-sample.py
```

» Disconnect power from the Pi, which should continue to run from the UPS battery. The time should be displayed every 10 seconds, and when the battery is below a certain level, you're told the Pi will shut down in two minutes (or immediately if you chose to skip the delay). Our Pi Zero with external Wi-Fi adapter ran for up to an hour before shutting down—ample time for getting the power back up.

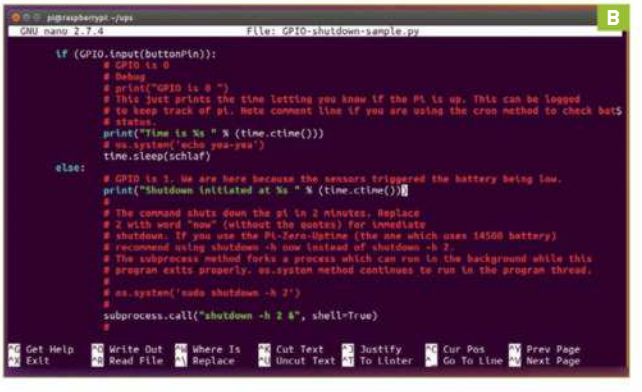

3 PERMANENT PROTECTION

The script must be run manually on every boot—for permanent automatic protection, reconnect your power, and move on to the final stage. The "GPIO-crontab" script is required to ensure the GPIO pins are tracked and a controlled shutdown is performed once the battery falls below the accepted threshold. Set it up thus:

```
$ sudo cp ~/ups/GPIO-crontab.txt /usr/local/bin/uptime-gpio.py
$ sudo chmod 755 /usr/local/bin/uptime-gpio.py
$ sudo crontab -e
```

» Choose your editor if prompted, then add the following line to have the script run silently once every two minutes (change `*/2` to `*/1` to run it every minute):

```
*/2 * * * * python /usr/local/bin/uptime-gpio.py
```

» Save the "cron" file, exit, and reboot. Now test the battery again—now, two minutes after the low-battery warning is reached, your Pi should automatically shut itself down. Congratulations, you've just ensured your Pi will never again suffer due to a sudden power outage!

Use Your Pi as a GPIO Expander Device

YOU'LL NEED THIS

RASPBERRY PI ZERO
The budget board costs just $5.

THE RASPBIAN OPERATING SYSTEM, based upon Debian, has long powered the Pi, and now it's available for desktops and laptops, reviving tired old computers with a light and usable operating system called Raspbian Desktop. In this tutorial, we're using Raspberry Pi Desktop on a spare computer, then attaching a Raspberry Pi Zero as a GPIO expander device using the micro-USB port of the Zero in USB OTG (On The Go) mode as a USB gadget. This gives our laptop the same OS as our Pi, and grants it access to the GPIO for projects. **–LES POUNDER**

1 GETTING STARTED

Downloading the Raspberry Pi Desktop is easy. Visit www.raspberrypi.org/downloads/raspberry-pi-desktop and click the "Download" link. This downloads an ISO file—a disk image that we can write to a spare USB flash drive. The best and easiest tool to use to write the image is Etcher (https://etcher.io).

›› Download Etcher and extract the file to your desktop. Now go to your desktop and double-click the file to open Etcher. The program asks you to tell it where the file is that you wish to write (ours is the Raspberry Pi Desktop ISO). Insert your USB flash drive—Etcher detects the drive and uses it. The final step is to click "Flash" and supply your password to write the image to the flash drive.

›› Once the flash drive has been created, remove it from your computer, insert it into an older machine, and power up. Depending on the machine, you'll need to select the boot device, either via the BIOS or by pressing a key when booting. On first boot, choose to run Raspberry Pi Desktop from the USB flash drive, and use persistence to save configurations and files to the drive without touching the hard drive inside the computer. The boot process should take around a minute. When completed, you'll see the Raspbian desktop.

2 GPIO EXPANDER

Something new that comes with Raspberry Pi Desktop is GPIO Expander. This enables a spare Raspberry Pi Zero to be used as a slave device, offering its GPIO for use with Scratch 2 and Python. It uses a USB/Ethernet gadget mode for the Pi Zero, which makes it possible for it to be configured as an Ethernet-enabled computer running a remote GPIO pin service, and means we can control the pins over a network connection—in this case, a USB connection. To use it, open a Terminal in the desktop, and type:

```
$ sudo apt update && sudo apt install usbbootgui
```

›› After a few moments, all of the necessary software will be installed. Now take your Pi Zero and a micro-USB lead. Insert the micro-USB connection into the micro-USB port marked "USB" on the Pi Zero. Insert the larger USB plug into a port on the Raspberry Pi Desktop computer. A few seconds later, a menu appears, asking you to select the role for the device [**Image A**]. Choose "GPIO expansion board," and click "OK." A new dialog appears and goes through the steps to set up the Pi Zero with the necessary configuration.

›› We now have only a couple more steps to take so that we can use GPIO Expander with Python 3. The first step is to export the GPIO Zero pin factory, so that we can use Pigpio for working with the GPIO. This is completely transparent to the end user.

›› In a Terminal, type the following:

```
export GPIOZERO_PIN_FACTORY=pigpio
```

›› The second and final step is to tell Pigpio where to find the GPIO pins, and this is via our Pi Zero, connected as an Ethernet device (over USB). Here we use an IPv6 address to identify the Pi.

›› In a Terminal type the following.

```
export PIGPIO_ADDR=fe80::1%usb0
```

3 PYTHON 3 PROJECT

Let's use our new GPIO Expander to create a quick test project using GPIO Zero to create a random sparkle LED project. We press the button and the Python code generates a random number used to blink the single RGB (red, green, blue) LED. The wiring for this project can be seen in the diagram opposite.

›› To code the project, we'll use the Python 3 editor, found in the "Programing" menu. Click "File → New" to create a new file, and we'll write the code in there.

›› First, we import from four libraries [**Image B**]. From GPIO Zero, we import the LED and Button classes, and from Time we import the sleep function used to pace our code. Then we import the uniform function from Random, which is used to generate a random float value. Finally, we import pause from Signal, to prevent our code from ending.

```
from gpiozero import LED, Button
from time import sleep
from random import uniform
from signal import pause
```

›› Then we create four variables used to identify where each pin of our RGB LED and the button are connected to the GPIO.

```
red = LED(17)
blue = LED(27)
green = LED(22)
button = Button(10)
```

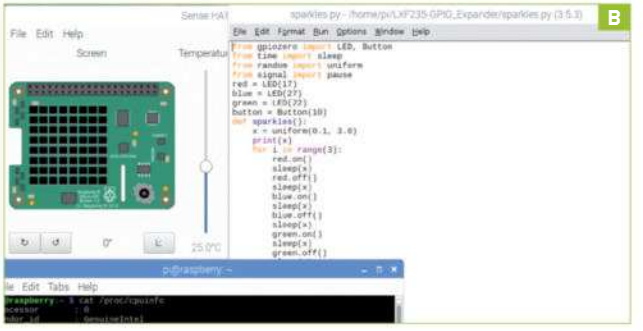

» In order to control our RGB LED [Image C], we are using a function called `sparkles`. This function chooses a random float value between 0.1 and 3.0, which is then printed to the shell for debugging purposes:

```
def sparkles():
    x = uniform(0.1, 3.0)
    print(x)
```

» We next use a "for" loop that will iterate over a sequence three times. The code inside the "for" loop turns the red, blue, and green LED on and off, using sleep to create a blink effect, thanks to x, our randomly chosen float value:

```
  for i in range(3):
      red.on()
      sleep(x)
      red.off()
      sleep(x)
      blue.on()
      sleep(x)
      blue.off()
      sleep(x)
      green.on()
      sleep(x)
      green.off()
      sleep(x)
```

» Our last two lines of code are used to handle a button press. Here it waits to run the `sparkles` function when the button is pressed. Then we use "pause" to prevent the code from exiting.

```
button.when_pressed = sparkles
pause()
```

» Save the code as "sparkles.py" to your home directory (/home/pi/), and exit the Python 3 editor. Now open a Terminal, and type the following to launch the code:

```
$ python3 sparkles.py
```

Press the button, and watch the RGB LED light up!

» We have successfully re-used an old computer and a spare Pi Zero, with which we can now create physical projects using many different languages.

HOW DO RASPBERRY PI DESKTOP AND RASPBIAN DIFFER?

We're using two new Pi-based technologies here. The first is a live distro based on Debian 9, offering Raspberry Pi Desktop for x86 computers, while the second is an implementation of remote GPIO control over a network. Both are handy if you have a spare computer and a Raspberry Pi.

The Raspberry Pi Desktop experience has a few noticeable differences from that found on the Pi. First, there's no *Minecraft*, due to licensing. In the Raspberry Pi Configuration tool (the GUI and raspi-config), the "Interfaces" section only has SSH, which means no I2C, SPI, Camera, and so on. This also means that any add-on boards that use these interfaces won't work—for example, Pimoroni's Explorer HAT, 4Tronix's Picon Zero, and Pi Supply's PiJuice. Boards that use the GPIO directly, either as a digital IO or using PWM, should work as expected, but testing is required before basing a project on it.

Right now, Raspbian Desktop is more about the look and feel of the desktop experience. Our GPIO project uses remotely controlled pins, with code being written on the computer, then executed on the Pi Zero. Check out the Raspberry Pi Foundation's Ben Nuttall's guide on his blog: http://bennuttall.com/raspberry-pi-zero-gpio-expander.

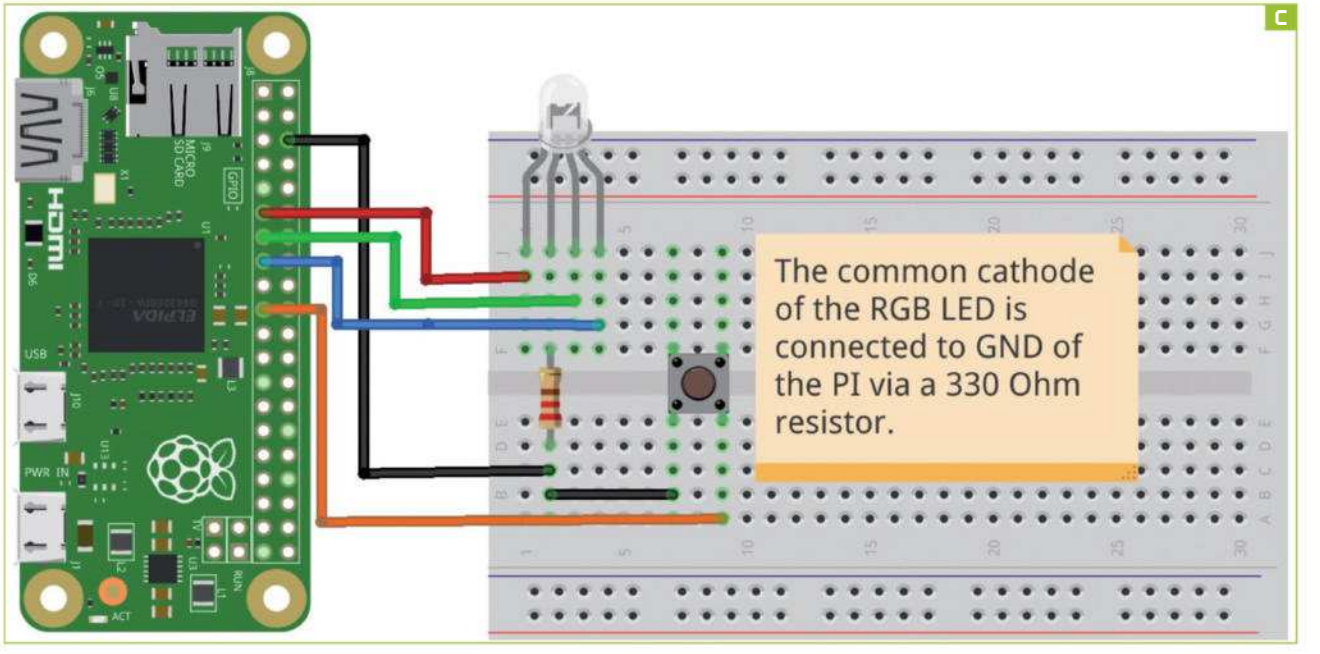

Make a Pi File Server and Ad Blocker

YOU'LL NEED THIS

RASPBERRY PI 3 OR 3 B+
With an Internet connection, USB flash drive or hard drive, keyboard, mouse, screen, and 2.5A official power supply, plus a PC running Linux or Windows.

TO CELEBRATE THE RELEASE of the new Raspberry Pi 3 B+, you're going to use it to create your own Samba server to store files and media for access across your home network. You'll then use the same Pi to block ads, and create a better browsing experience.

You first need to make sure that you have the latest version of Raspbian, so download this from the Raspberry Pi website, then using your favorite tool, flash the image to a blank microSD card. When completed, eject the card, and place it into your Pi, then connect your keyboard, mouse, screen, Ethernet, and finally the power. Your Raspberry Pi boots up, and because this is the first boot, it expands the filesystem to fully utilize the space on your SD card. –LES POUNDER

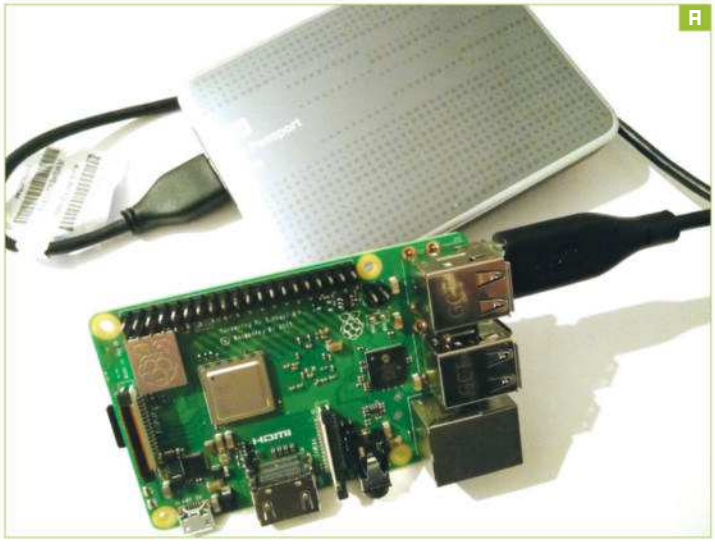
A

1 GETTING STARTED

Soon you'll see the Raspbian desktop, so now update your Raspberry Pi 3 B+ [Image A]. Open a terminal; the icon is in the top-left of the screen. In the terminal, type the following, line by line, and press Enter at the end of each line:

```
$ sudo apt update
$ sudo apt upgrade -y
```

» To ensure that you can always find the Pi on the network, let's fix its IP address. Right-click the networking icon, and select "Wireless and Wired Network Settings." In the "Network Preferences" dialog, configure interface eth0. Select "Automatically configure empty options," then change the IPv4 address to match that of your network. For example, our network uses 192.168.0.xxx, so we chose to fix the IP address to 192.168.0.100. Click "Apply," then "Close." Now reboot the Pi to ensure that the IP address has been changed and to check you can open a terminal, and type:

```
$ hostname -I
```

» With the terminal still open, let's install the Samba package. Type the following, and press Enter:

```
$ sudo apt install samba samba-common-bin
```

» With everything installed, you now need to insert your USB storage device, which you'll use as your shared storage. Raspbian automatically mounts the drive for you, but you need to know where it's mounted. In the terminal, enter the following command:

```
$ mount | grep /dev/sd
```

2 SET UP SHARED STORAGE

If there's only one USB stick inserted, you should see a device labeled "/dev/sda1." This is your USB stick, so make a note of the device, and where it's mounted, typically in /media/pi/. Make a note of the full path to the drive. You now have the device that's your USB drive, and it has a directory that you wish to share using Samba. Let's tell Samba where to find the drive. In the terminal, type the following to edit the smb.conf file:

```
$ sudo nano /etc/samba/smb.conf
```

» When the file loads, go to the bottom of the file, and recreate the following lines:

```
[share]
Comment = Pi shared folder
Path =  /media/pi/<NAME OF DRIVE>
Browseable = yes
Writeable = Yes
only guest = no
create mask = 0777
directory mask = 0777
Public = yes
Guest ok = yes
```

» To save and exit the editor, press Ctrl-O, then Enter, then Ctrl-X. Now let's ensure that the default user, pi, is added as a Samba user. In the terminal, type:

```
$ sudo smbpasswd -a pi
```

» Next, to ensure that your Samba config is correct and ready to use, you need to tell Samba to restart, so in the terminal, type the following:

```
$ sudo /etc/init.d/samba restart
```

3 TEST THE SERVER

With the installation complete, you now move to your Windows or Linux computer to test that you can find and use the Samba server. Samba is cross-platform, so Windows, Mac, Linux, and mobile users can access one central location for their files [Image B].

» On Ubuntu, you can use the file manager (Nautilus) to connect via the "Network" shortcut. The name of your server "RASPBERRYPI" should be visible. Click "share," then enter the username "pi" and password "raspberry" to access the server. The process is similar for Windows 10 devices, because the "Network" shortcut is also found in File Manager. Android users can also download Samba client apps from the app store, and access files just like their desktop counterparts. Now you have a single point where

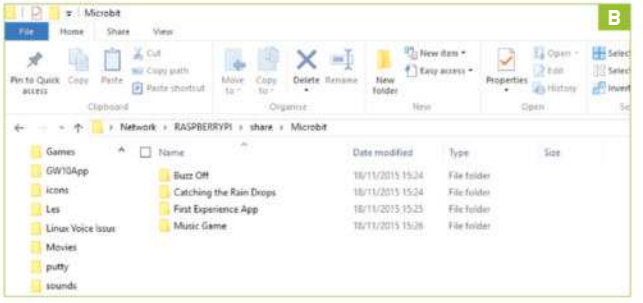

you can save files, and share them across the network—handy for the large amounts of media consumed by busy households!

4 SET UP PI-HOLE

We all dislike advertisements on web pages, but installing an ad blocker on every device is time-consuming. What if you created one device that filtered all of this for you? Pi-Hole [Image C] is just that device. Running on your Raspberry Pi, Pi-Hole replaces your standard DNS server (normally your ISP, unless you change it) with its own that filters out the ads. Oh, and it's very easy to install, so let's get started.

›› You need to open a terminal, and type the following command to download and run the install script:

```
$ curl -sSL https://install.pi-hole.net | bash
```

›› Note that installation via this method carries a risk if being used with unknown sources. If in doubt, download the source code, and examine it before use.

›› Pi-Hole installation is an automated process, and requires a minimum amount of input. But there are times where you're asked questions. One of these questions asks you to set up a static IP address—this is unnecessary, because you already have a static IP created for the Samba server (if you haven't set up a Samba server, you can set the static IP address using Pi-Hole, but make a note of the IP address).

›› You're also asked who your upstream DNS provider is. You can choose from a list: Google or OpenDNS are popular choices. Follow the rest of the installation steps, and when asked to install the web admin, select "On." After a few minutes, installation is complete, and you need to power down your Pi, and remove all of the connections except for the Ethernet and power lead.

5 INITIATE AD BLOCKING

Power up the Pi, and wait for it to boot—it takes around 30 seconds. On your Linux/Windows computer, go to your network settings. On Ubuntu, you can right-click the "Network" icon, select "Edit Connections," click the interface that you wish to edit ("Ethernet" or "Wi-Fi"), then click "Edit." In "IPv4 Settings," change the DNS server to match the IP address of the Raspberry Pi. Click "Save," then "Close," then click the "Network" icon, and click "Disconnect" for the interface you've just configured. Reconnect to force the interface to connect using the Pi-Hole DNS server.

›› You now have a transparent ad blocker on your network. It works across all devices, including Android, where you can specify a static IP address and DNS server by a long press on the Wi-Fi AP name. Next, select "Modify Network → Advanced Options," then enter the desired static IP and DNS server (your Pi Hole device) address.

›› If you love stats, Pi-Hole has plenty. If you open a browser to the IP address of the server followed by "/admin" you can see a dashboard interface that logs all of the queries blocked for your device. If you log in using the password provided during the install, you can see even more details.

›› So there you have it—two services using the power of the new Raspberry Pi 3 B+'s enhanced networking to keep your browser sessions free from ads, and your files saved in a convenient location.

ALTERNATIVE APPROACHES

Samba is easy to set up, and it works across operating systems. It can also be used to share printers over a network using the CUPS service. CUPS is easy to install—you can find plenty of guides online.

NFS also offers benefits for Linux users. Chief of these is that a user's Home directory can be served over the network via NFS—handy for thin client devices, while storage devices, such as optical drives, can be accessed as though the user is sitting at that remote machine. Installing NFS is simple, and there's a guide at https://help.ubuntu.com/community/SettingUpNFSHowTo that works with Raspbian.

If all of this file server talk has sparked an interest, we recommend using Diet Pi (http://dietpi.com). It comes with an automatic install script for many different file servers, Pi-Hole, and other servers for home automation, media cataloging, and torrents. All of which can be powered by the humble Raspberry Pi—pretty impressive.

Check Speed After Vulnerability Fixes

YOU'LL NEED THIS

INSPECTRE
Simple program that toggles Spectre and Meltdown protection on/off. Get it from www.grc.com/inspectre.htm.

X264 BENCHMARK V5.0
For stress-testing your CPU via 1080p video encoding. Download it from https://archive.techarp.com/showarticlea96b.html.

THEY SOUND LIKE BOND MOVIES—actually, one of them was a Bond movie—but in a computing context, Meltdown and Spectre are nothing like as much fun. The pair of exposed and fearfully publicized CPU vulnerabilities caused much gnashing of teeth earlier in 2018, due to the potential for hackers to target two decades' worth of processors. Teeth were gnashed even harder when the patches that Intel and AMD rolled out to fix these vulnerabilities were reported to decrease system performance.

The extent of that performance hit isn't certain, though, and varies depending on your CPU model, your OS, your storage devices, and several other factors—which is why it's time to roll up the metaphorical sleeves, and do some testing. Using a tiny app that toggles Meltdown and Spectre protection on and off by enabling or disabling their associated patches, and running some CPU benches in both environments, you'll get an understanding of the hit your system took to keep you safe. Note that we don't recommend leaving the patches disabled simply to restore that lost performance—what's a few frames per second when your personal data is up for grabs, really? This is purely an exercise in assessing the damage—gulp. –PHIL IWANIUK

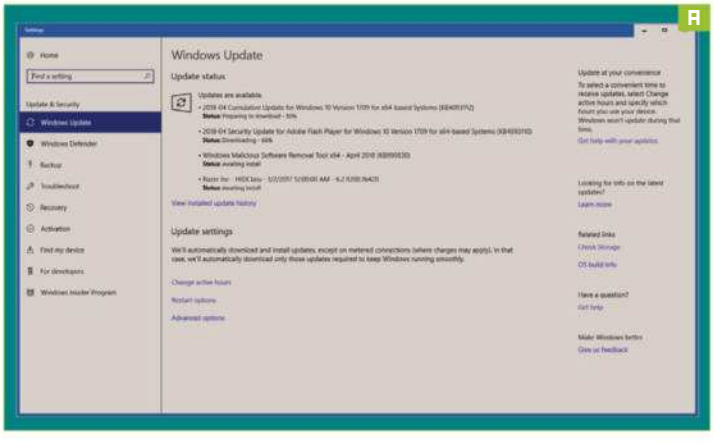

1 ALLOW ALL MELTDOWN AND SPECTRE UPDATES

As obvious as this sounds, there's no point running CPU benchmarks to check how your performance has been affected by patches you may not even have installed yet. So, let's make sure you have. Probably by now, your OS, browser, and other programs have done their business, but just in case: Navigate to the Windows Update by typing "Update" in Cortana, and clicking the program when it appears. Under "Update status," click "Check now" to ensure you don't have any updates pending. If you do [Image A], get them installing right away. It's also worth checking that you're running the latest version of your web browser, and enabling any program updates that prompt you to install them. That includes antivirus software.

2 INSTALL THE REQUIRED SOFTWARE

You need two programs to properly assess your before and after performance in relation to the security patches you just installed. The first, InSpectre, immediately tells you whether your system's protected against the two exploits. We're using the x264 1080p encoding benchmark [Image B] to assess performance after that, but you're obviously free to use any CPU benchmark you like, as long as it spits out a number at the end that you can use to contrast and compare with other runs. If InSpectre tells you your system isn't currently protected at this point, simply repeat step 1. Do not pass "Go," do not collect $200. Otherwise, with this and your benchmarking tool installed, you're ready to progress.

3 PREVENT INTERRUPTIONS BY GOING OFFLINE

The goal here is to run the same benchmark on the same system, under exactly the same conditions. The only variable should be whether or not your CPU security patches are enabled. In order to do so, let's eliminate the potential for Internet connection or storage drive activity to mess with the results. You never know when GeForce Experience might decide now is a good time to install new graphics drivers and plunge your screen into temporary darkness, after all. First, go offline temporarily, whether that means pulling out an Ethernet cable or USB wireless receiver, or simply disabling Wi-Fi at the bottom-right of the Start bar [Image C].

4 GET YOUR STORAGE DRIVE FIGHTING FIT

Meltdown and Spectre's fixes don't just impact CPU performance in a vacuum. Your CPU is in charge of software-defined storage (such as partitions), and the size of that workload is, most agree, affected by the security

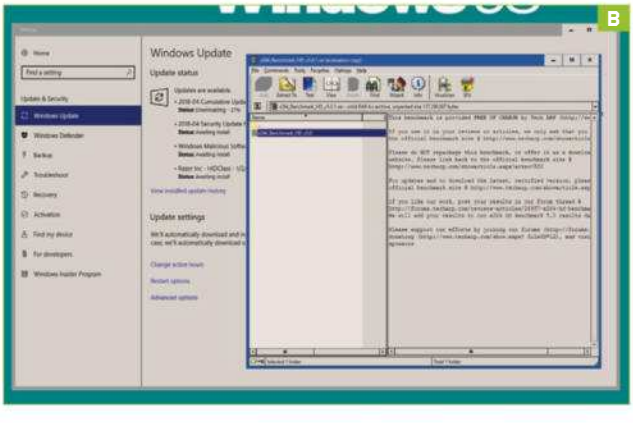

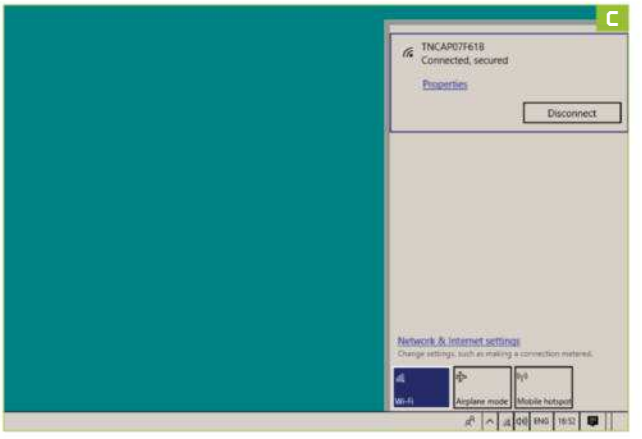

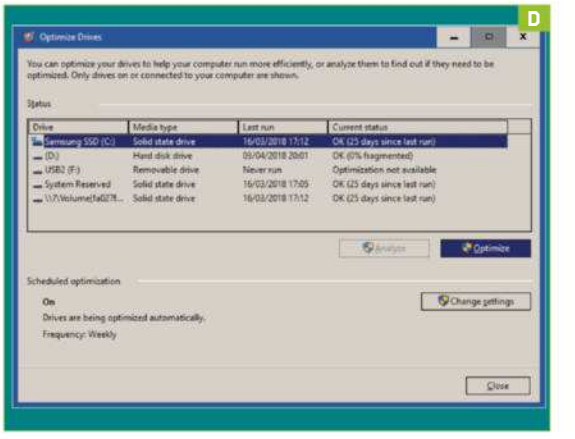

patches. All of that means that if we want our testing to be fair and accurate, we need disk activity to be the same during both test runs. That's why we've gone offline, and it's why we recommend running a TRIM on your SSD before heading to the benchmarks. Don't sweat this if Windows is installed on to an HDD on your machine, because the numbers are probably so slow in the first place that performance differences will be negligible. SSD users should right-click their primary drive in Windows Explorer, then head to "Properties → Tools → Optimize" [**Image D**]. This starts the TRIM—it's essentially a defrag for solid-state drives, clearing up data cells, and tidying up where data is written. Now when you run the benchmark, you'll know that your disk isn't installing something in the background that it just downloaded online, or silted up by not having run a TRIM command for too long.

5 BENCHMARK, THEN BENCHMARK AGAIN

Now's the time to actually collect the data you can use to analyze any performance hit. You can do the following procedure in any order, but for simplicity's sake, we'll run the first benchmark with Meltdown and Spectre patches disabled. Open InSpectre to do this, and click that ominous "Disable Meltdown Protection" button in the bottom-left of the program. Then, having summoned the courage once more, click the "Disable Spectre Protection" button next to it. Good job we went offline, huh? InSpectre should now tell you that you're not protected in scary red letters. Restart your system, making sure you're still offline, and wait for any startup programs to get up and running. Now it's time to run your benchmark. We're taking three 1080p encoding runs in x264 to work out an average, just for a little more accuracy down the line.

›› With those three numbers gathered and logged, we'll grab another useful bit of data by unRAR-ing the x264 install folder on to our desktop three times, and noting down how long it takes each time. This gives us an insight into disk access, as well as raw CPU grunt in a synthetic benchmark.

›› It's finally time to wrap ourselves back up in the warm blanket of Meltdown and Spectre protection, so go ahead and open InSpectre, and enable the two options on the lower part of the window. Reassuring green writing should now tell you that you're safe again. Phew! Give your PC a restart, re-enable your Internet connection, and once again allow all startup programs to get back up to speed. Finally, run x264 or your chosen benchmark three more times, then unRAR the installer folder on to your desktop three more times, and gather the data.

6 OBSERVE THE DIFFERENCE, WEEP INTO HANDS

The differences in your numbers will vary from system to system. Probably, they represent a 3–15 percent decrease in performance with Meltdown and Spectre protection enabled, with your decompressing numbers falling toward the higher end, and CPU benchmark performance varying more marginally. If you're seeing differentials much bigger than that, there's almost certainly another factor at play, so don't go cursing those fixes just yet. Simply run more tests, ensuring other variables aren't messing with your figures.

›› Is a 3–15 percent performance hit acceptable? Considering the dire consequences of someone using the Meltdown or Spectre exploits to access your passwords and other personal info from your kernels, we think it is. No shrinkage in system grunt is to be welcomed with open arms, but thankfully it looks like the impact of these fixes is unlikely to prove noticeable in everyday usage.

BOILING POINT

How did these two notorious CPU problems come to be in the first place? And why did it take their manufacturers over 20 years to realize their architectures were at risk? The answers to those very pressing questions still haven't been provided, and perhaps even more chilling is the fact that they were discovered by complete accident by three different research teams, working independently, at almost exactly the same time. That the exploits affect CPUs dating back to 1995 but weren't spotted and shown to Intel until June 2017 suggests that many previous researchers were aware of it, but didn't act upon it. It's more comforting to imagine that it was blind coincidence, though, so let's ensconce ourselves in that theory.

Prior to their discovery, it was long held true that virtual memory data held in the CPU cache was as good as under lock and key. Processes simply couldn't access that data without the proper authorization—so everyone erroneously thought. Both Meltdown and Spectre work in a broadly similar way, accessing the CPU cache and all the riches it contains. And because the problem is fundamental to CPU architecture that has been iterated on for over 25 years, nearly every Intel and AMD processor produced during that period is affected.

Enhance Your Windows Desktop

YOU'LL NEED THIS

WINDOWS 10
And the desire to change its desktop to something better.

YOU WIMP! Readers as old as our executive editor, Alan, will know that Windows Icons Menus and mouse Pointer (WIMP) is the ancient acronym that describes the original mouse-driven graphical interfaces that span out of Xerox labs. Despite Windows 8, Microsoft was never able to kill off the classic icon-based desktop that most PC users know and love.

We should admit to the truth, though: The classic desktop is terrible; it's as bad as a real desktop. You'll leave your garbage scattered all over it—empty coffee mugs, Oreo crumbs, and worse. It is, however, a little draconian to utterly ban the ability to keep recent items left handily on your desktop. Surely there's a helpful middle ground?

The smart news is there's a host of clever, useful, and possibly silly add-ons for Windows that push the desktop to the next level of functionality. We're going to quickly run through what we think are the best five(ish)—yes, it's a listical—that are worth your attention. –CONNOR MACLEOD

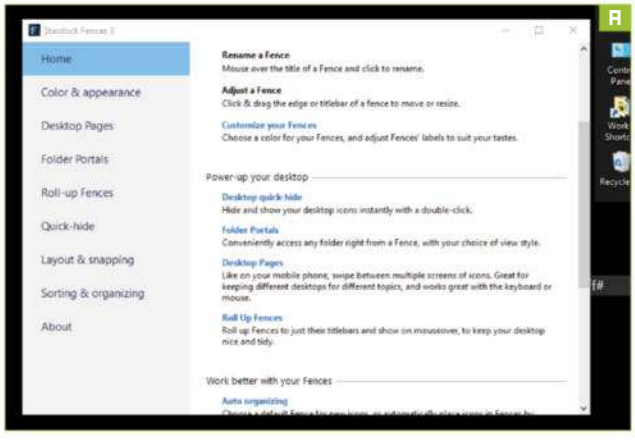

1 SITTING ON THE FENCE

A shorthand answer is to just buy the entire StarDock Object Desktop suite from www.stardock.com. For a cent shy of $30, it's a lot of slick software for your dollar. For us, though, the star of the show is Fences [**Image A**], for which there's no real alternative anywhere. Fences automatically groups icons on your desktop into flexible, movable boxes. These can be moved, hidden, and resized on the fly, while a genius part is they'll all automatically scale with your desktop resolution, and as monitors are added or removed, which includes their size and positioning. It's worth its weight in gold, and the smart automatic setup can instantly tidy up the most messy of desktops. It also retains the standard desktop, so that's still there below your fences.

›› Did we just say there's no real alternative to Fences? We're not aware of anything as slick, but a freeware alternative called Nimi Places (http://mynimi.net) has been around for years, and visually looks good, but takes more work to set up than Fences, and we find the way it handles icons a little odd.

2 DON'T GO CHANGING

We don't like change around here.... Do you remember Windows 8? Those were crazy days, right? But you need to upgrade Windows to keep on top of security updates, and get access to the latest releases of DirectX, and the rest. Thankfully, there's an easy solution in the form of ClassicShell (www.classicshell.net), where there's good news and some bad news.

›› The good news is that it remains the best way of transforming Windows 10 back to the classic-looking Windows 7 Start menu and theme you know and love [**Image B**]. The bad news is that the author ceased supporting it at the end of 2017, but they at least had the good sense to open-source the code, so fingers crossed, some hero will pick up support for this, and it can continue down the line. It's no joke that many people find their productivity hamstrung by the terrible changes to the Windows Start menu and other areas, so there's a genuine argument for changing things back, if that helps you.

3 TWEAK YOUR TASKBAR

The Taskbar is an unsung hero of productivity. It lists and enables you to control your running programs, and with Windows 10, it extends features to virtual desktops, too. Starting here and moving beyond, we stumbled across the awesome Winaero (www.winaero.com), which is a complete Taskbar, desktop, menu, Windows, and more tweaking tool for Windows 7, 8, and 10 [**Image C**]. We are trying to concentrate on actual productivity enhancers, but it's good to see that Winaero covers off tweaking the color scheme of the default theme and the lock screen images.

›› Of more use are the options to control the Aero Shake and Aero Snap features. Dump looking in the Windows

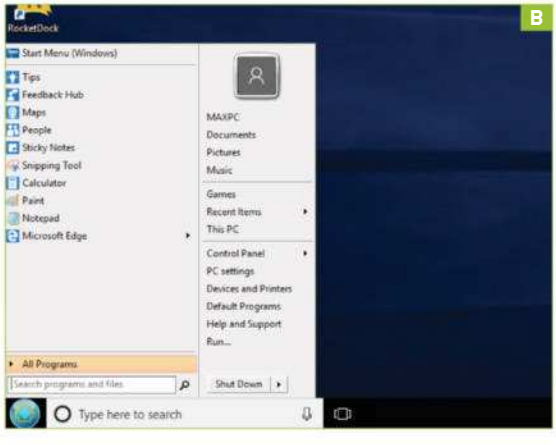

DOCKS

If you're not content with fixing the Windows Taskbar, why not replace it? Docks used to be all the rage when OS X appeared in 2000, but the clamor has died down more recently, replaced by the full-screen touch Launcher, we guess. It's a shame docks were often just visual fluff, as to truly replace the Taskbar, they need to offer task management, and more. Two well-known docks are Rocket Dock (https://rocketdock.com) and Circle Dock (http://circledock.wikidot.com/start), but both have been put out to pasture for almost a decade, though they still both work.

One of the few docks still maintained worth a look is WinStep Nexus (www.winstep.net/nexus.asp), which offers a fully featured free version and a paid-for extended option. It supports Win XP through 10. The paid-for version does enhance its features greatly, including nested folders, desktop replacement with icon management, monitoring widgets, themes, and more.

Store, and control Windows Updates more easily. Bring back the old volume controls, disable web searches, and kill Live Tiles in one fell swoop—we'd pay money for that feature. Customize the context menus, file explorer presets, and default tools.

›› Another easy and free way to get more from the Taskbar itself is 7+ Taskbar Tweaker (http://rammichael.com/7-taskbar-tweaker), which enables you to configure aspects of the Windows 10 Taskbar you otherwise couldn't. Turn off the grouping of pinned app windows, cycle between Taskbar buttons with your mouse wheel, assign a "close app" action to your middle mouse button, and much more.

4 FLUFFY FLUFF

We call FolderMarker (http://foldermarker.com) fluff, but in fact it's a genuinely useful tool. If you have a folder full of, erm, folders, how are you supposed to visually tell them apart at a glance? In Windows, you can't. Sure, system folders—"Pictures," "Downloads"—get their own icons, but user folders are all plain-Jane yellow. FolderMarker [**Image D**] offers a host of alternatives, not only plain colors but also ones with handy attached icons. These can then be easily applied via the right-click context menu.

›› Our other fluff-based recommendation is more of a throwaway, but it seems popular, and it's CustomizerGod (www.door2windows.com/customizergod), and the one job it does is enable you to change any and all system icons. We personally found it a little overwhelming, and would be a mammoth timesink for little gain, but we're mentioning for completeness.

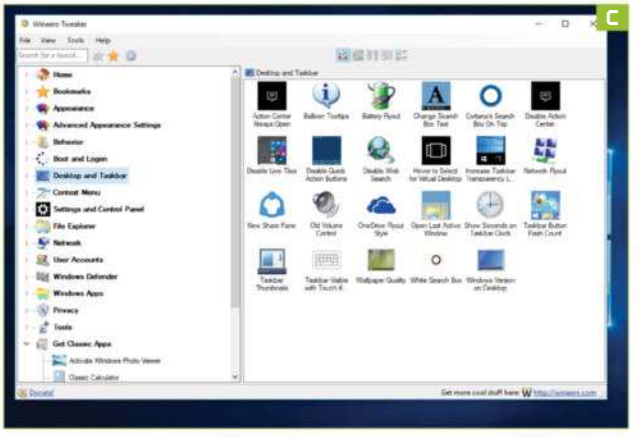
C

5 WINDING DOWN

Windows 10 does now include a night light feature as part of the Display settings, but it's rather limited in options, and can be unavailable depending on your display driver. In the past, f.lux (https://justgetflux.com) would be the go-to screen dimming tool, but we're going to suggest SunsetScreen (www.skytopia.com/software/sunsetscreen), which has a portable install, and a more flexible dimming schedule to match your body clock. So, now you can work to the wee hours of the night, and not destroy your natural biorhythms with blue-soaked light.

›› People can get very territorial about changes to their beloved desktops, but who can blame them? You spend years getting used to working one way, and almost overnight, you're expected to change to a new system. So many of the tools here are geared to changing new systems to work like the old ones, but choice is good, and you're free to choose what you'd like to customize or not!

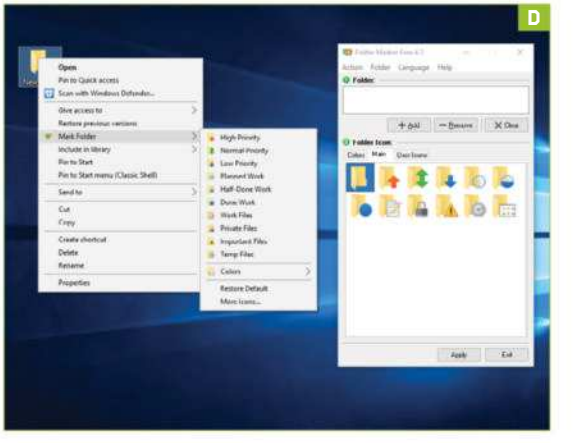
D

Create a Split-Toning Effect in Photoshop

YOU'LL NEED THIS

ADOBE PHOTOSHOP OR LIGHTROOM
Subscribe to a suitable package at www.adobe.com.

PHOTOGRAPHIC TONING once meant changing the color of a black and white image by replacing the silver used in the development process with a silver compound. This is why old photos have a wonderful sepia tone—it's down to chemistry, specifically silver sulfide in the case of sepia, and it had the side effect of increasing the lifespan of the resulting print.

Of course, this is relatively straightforward to do with digital post processing, and we can take the process a step further by toning the light parts of the image in one color, and the dark parts in another. This is known as split toning, and can be achieved in both Photoshop and Lightroom—the latter even has a panel dedicated to it.

It's not just a sepia effect you can produce. Selenium toning can add a lovely purple hue to an image, while compounds of other metals produce blues or reds. You don't need to worry about the chemistry, however, and can tone an image in whatever way you think looks best, even if it's never been produced in any darkroom before. –**IAN EVENDEN**

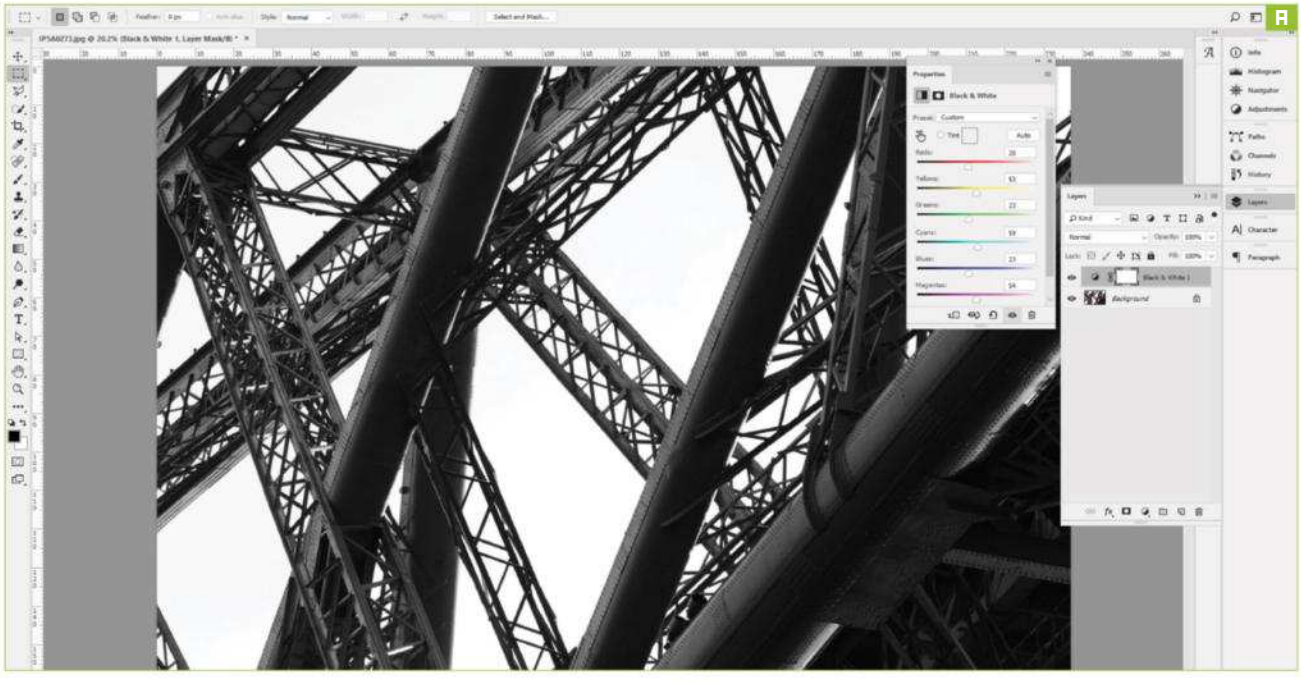

1 IN BLACK AND WHITE

First of all, you need to make your image black and white. It's not enough to make it grayscale in Photoshop, however, because we need to put color information back in. There's a couple of ways you can achieve this: by changing the color mode ("Image → Mode") to grayscale and immediately back to RGB, or by placing an Adjustment Layer above your background layer, using the button that looks like a circle divided between black and white halves at the bottom of the Layers palette. You can use a Saturation Adjustment Layer, and drag the slider all the way to the left, or use the specialized Black and White layer [**Image A**]. This offers a load of options in its "Properties" window for how colors in your image are converted to shades of gray, and also an "Auto" button for those who don't want to fiddle too much. We're not going to leave the image in grayscale anyway. Now's a good time to save as a PSD, so your new layer structure is preserved.

2 HIGHLIGHTS OR SHADOWS?

Place another Adjustment Layer on top of the one that created the grayscale effect. Choose "Color Balance" from the "Adjustment Layer" menu, and it should appear at the top of the layers stack—if it doesn't, drag it there with the mouse. It opens a "Properties" window; in here, you can mix the colors that will decide your final tone. Choose which part of the image to affect—highlights or shadows—from the drop-down menu at the top.

3 WARM OR COOL?

The pairs of colors you choose can greatly affect the mood of the image. Warm colors advance toward the viewer, while cool colors recede into the background. Apparently, this is something to do with where different colors come to focus within the structure of the eye, but this

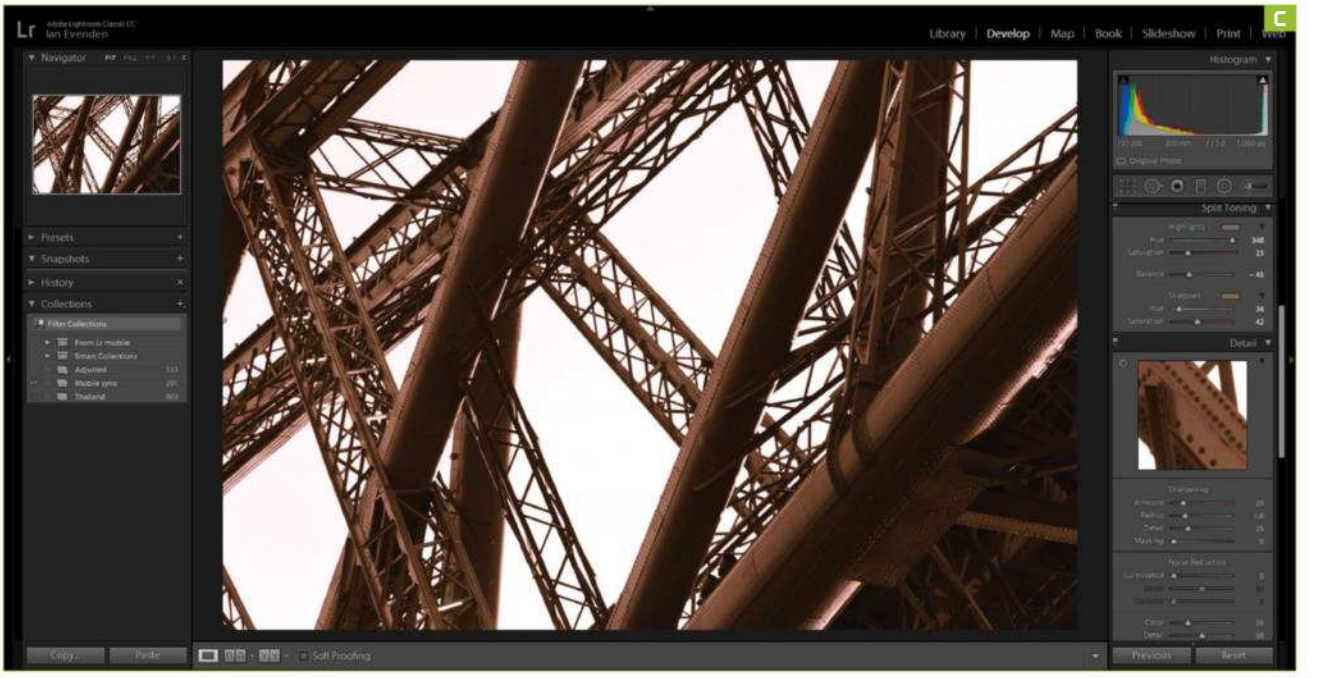

is *Maximum PC,* not *Ophthalmology Monthly,* so just take our word for it. Other color schemes to try include complementary colors—those that are next to one another on the color wheel—and triadic colors, where you draw a triangle on the color wheel, and combine the colors marked by its points.

4 OTHER METHODS

There are other ways of achieving the same effect. Take your image back to a point after you created the Black and White Adjustment Layer, but before you created the Color Balance one, and create a Gradient Map Adjustment Layer, ensuring it's at the top of the stack. Drop down the menu next to the selected gradient, and you have access to all the others in the set. Experiment with these, but notice there's a cogwheel that can be clicked—this drops down yet another menu, from which you can choose further sets of gradients, including Photographic Toning .

5 GORGEOUS GRADIENTS

There's a lot you can do with these gradients, including creating custom ones. Click on a gradient in the "Properties" window, and you open a "Gradient Editor" [**Image B**] that lets you choose not only the colors involved, but when they begin to blend, and how much of one color is preserved before it blends into the next. There is also a "Reverse" checkbox, which can be used to create some starkly beautiful results.

6 LIGHTROOM OPTIONS

Then there's Lightroom, where you can apply the same effect to a Raw image, export it, then change the settings completely and easily. Open your image, and head to the "Develop" module, where you can reduce the Saturation to get a grayscale image using the slider. Head down the list of modules, and you'll find one dedicated to split toning [**Image C**]. It's a bit bare-bones, but includes a "Balance" slider that affects the balance between the two colors you've chosen. Pick your highlight and shadow colors using the "Hue" sliders, then their strength using the "Saturation" sliders, and blend them with "Balance." It's one of the easiest ways of producing the split-tone effect, and because of the way Lightroom works, retaining your Raw file as a "negative," and creating new files each time you export, you can produce endless variations.

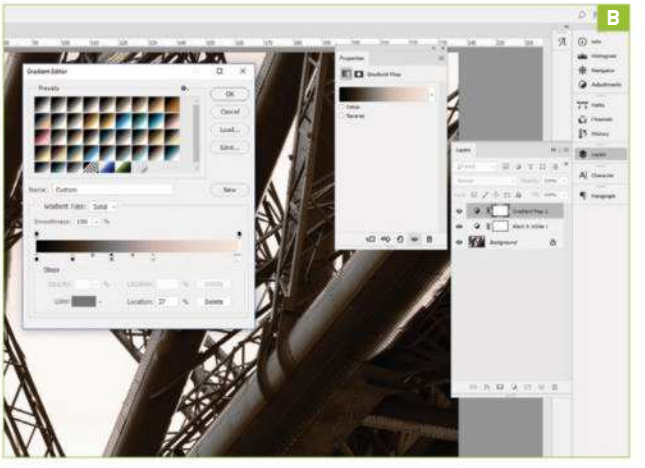

COMPLEMENTARY COLORS

Split toning is effective on different types of photo. Architectural shots, like the iron bridge we've used here, take it very well, their sharply defined lines and overlapping parts allowing plenty of areas where shadows and highlights intersect. But it's also an interesting thing to try with portraits, because its ability to inject mood into an image can often be used to reinforce the photographer's intent when they shot their subject. Adobe has a web app—at http://color.adobe.com—that can help with choosing colors, spitting out RGB values for complementary, analogous, and compound colors, and a great many more types besides. You choose a base color, then the type of colors you'd like to generate, and the app does the rest.

BUILD IT

ZAK STOREY, DEPUTY EDITOR

The Zenith of High-End Builds?

AMD's 12nm manufacturing process makes its debut in our Build It section

LENGTH OF TIME: **2–3 HOURS**

LEVEL OF DIFFICULTY: **EASY**

THE CONCEPT

IT'S GOOD to have some competition back in the industry once again. Seriously, how long were we plodding on with on mainstream quad-core processors? Over eight years? Staggering. The reality is that the computing industry thrives on competition, and as years pass, and more and more manufacturers hit the dirt, the likelihood of us ever seeing such a large paradigm shift again—certainly when it comes to processor performance—is slim at best.

AMD's Zen architecture is without a doubt the catalyst for the retirement of that mainstream staple quad-core part. And with it we've seen an accelerated processor arms race between the two top chip manufacturers. As Intel scrambles to try to reclaim the consumer multithreaded performance title, with six-core parts on sale now, and eight-core parts rumored to be just around the corner, it does make you question what exactly it was that was holding Intel back all this time?

Alas, we digress—this isn't about who did what best, this is about the launch of a new processor generation: AMD's first foray into its Ryzen refresh. Its original lineup of Ryzen processors rocked the world with their multithreaded price to performance. But single-core? Well, that had problems, serious problems, and those were reflected in game, too. Now, AMD tells us, it has almost entirely remedied that situation, and with the disparity between Intel and AMD's latest parts nearing a meager 1 percent performance difference across a plethora of titles, is it time to jump ship to Team Red? We investigate.

SECOND-GEN JOY

SO, THIS SYSTEM is a little nutty—it doesn't exactly fall into the category of an efficient gaming machine, nor does it delve into the realms of enterprise workstation. In short, we were given one mission statement for this build: Craft a rig housing Zen+ that looks good. So, that's exactly what we aimed for. Couple that with this particular writer's high-end balanced approach to hardware selection, and what you're left with is a top-level machine that's perfectly optimized for both throwing down in-game frame rates, and ably finishing any editing task you have to hand.

As is typical of us here at *Maximum PC,* the storage we've chosen is, by far, the biggest expense in this build. That's no chump-change hard drive for the OS there, no sir. That's 480GB of the finest 3D Xpoint Intel Optane solid-state storage. On tap, ready to deliver premier-level performance for any and all editing applications out there. Unfortunately, it's not cheap. At $576 for the exotic plug-in PCIe storage stick, it's well over the $1 per GB mark ($1.20 to be exact). But, oh boy, does it rip up the playbook when it comes to random 4K read and writes—arguably the most important metric when it comes to SSD performance.

On top of that, we've thrown in a hefty 1TB Samsung 860 Evo for any additional storage needs, and some (still excessively expensive) Corsair Dominator Platinum DDR4 sticks, which clock in at a total of 32GB (4x 8GB, dual-channel), running at a fairly average 3,200MT/s clock. All of which is ideal for Ryzen's fabled infinibro fabric. With that out of the way, then, let's get on to the build.

INGREDIENTS

PART		STREET PRICE
Case	NZXT H700i	$180
Motherboard	MSI X470 Gaming M7 AC	$260
CPU	AMD Ryzen 7 2700X	$330
CPU Cooler	Corsair Hydro H115i Pro	$140
GPU	MSI GeForce GTX 1080 Gaming X 8G	$700
Memory	32GB (4x 8GB) Corsair Dominator Platinum @ 3,200	$500
PSU	850W Corsair RM850x White 80+ Gold	$170
Storage 1	480GB Intel Optane SSD 900P	$576
Storage 2	1TB Samsung 860 Evo 2.5-inch SSD	$280
OS	Windows 10 Home 64-bit	$100
Total		$3,236

1 H700I LAID BARE

NZXT'S H700I has long been a favorite in the office, mostly because of the premium interior this thing has. Interior? Yup, building in it is a dream; it's as versatile inside as it is pretty on the outside, with most of the panels easily removable. One thing that is a bit of a disappointment with this chassis is the front SSD mount. If that were just a solid panel, and you mounted the SSDs on the PSU cover, it would have made the build look far cleaner. You could even incorporate some of that super-secret-sauce LED lighting into the NZXT logo, as we saw with the Manta.

2 COOLER MOUNTING

ALTHOUGH CORSAIR'S CPU BLOCK may be rather divisive aesthetically, there's no denying how beautifully easy it is to install it on the chip. Simply replace the stock Intel bracket locking mechanism on the bottom by rotating it to the right, then attach the thumb screws and locking brackets, and mount directly on to the stock AMD mobo brackets. Simple. It's arguably one of the more secure fitting solutions we've seen on the market to date.

3 RADIATOR BRACKETS

ANOTHER NEAT THING baked into the H700i is this snazzy radiator bracket, situated in the top of the case. Secured down by four thumb screws, you can lift it directly out of the chassis, and pre-install your rad and fans outside of the case. Then all you need to do is slot it back into place, carefully thread your tubing and fan cables through, and you're good to go. It's worth noting that we did have a bit of a conflict here with our 140mm fans and the stock LED strip NZXT has baked into the top of the chassis, forcing us to remove the strip.

4 EASY CABLES

MODULAR POWER SUPPLIES are great, up to the point where they're installed in the case. It's a small tip, but the best thing to do is pre-attach all the cables to the PSU outside of the case first, then thread them through the PSU slot, before securing the power supply in place. The H700i comes with a PSU bracket that you pre-mount to the PSU, then slide it into place, and secure it down first.

5 LED FAILS, SORTA...

TRYING TO FIGURE OUT how exactly to light this case was a challenge. As always, the best kind of lighting is that which is hidden. Keep the emitters out of sight, and just throw the light where you want it to go. We tried throwing it in the bottom first, but that conflicted with the cable hide bar's mounting points, and was also blocked by the SSD. We tried throwing it up top, but it was in view. So, the only other option was up at the front of the case. It's visible with the side panel off, but otherwise it's kept out of sight fairly well. If we'd have gone with a 240mm AIO in the roof, we might have been able to hide another strip up there, which would have been ideal, if we're honest. However, with the 280mm, it was just too tight up there.

6 INTEL OPTANE POWAH!

AHH, THE BEAUTY OF A JOINT AMD-Intel partnership. Going with the $580 Optane PCIe SSD was a last-minute decision. In short, it was a toss-up between this and a Samsung 960 Pro, but in the end, we simply couldn't resist the Optane's random read and writes. There's also something strangely satisfying about having a PCIe add-in card SSD in a rig, especially when you're running a single-card build like this.

1 **These streaky cables have been driving us insane. Corsair does include some cable combs with the PSU, but they simply don't fit—at least, our sample's didn't.**

2 **By default, MSI's X470 Gaming M7 AC supports dual eight-pin EPS power to better balance CPU juice. However, you can run with a single eight-pin just fine.**

3 **Front airflow is powered by a set of stealthily installed Noctua NF-A14 PWM Chromax fans.**

4 **Not that you can see them from this angle, but removing the 3.5-inch hard drive cage below the PSU cover helped give us a bit more wiggle room in the rear for all the cables we needed.**

RYZEN SHINE

A SMATTERING OF LIGHTING, a few braided cables, sexy silver heat pipes, and some sleek tempered glass was all it took to raise this build to the heady heights of a premium PC. The hardware within is powerful, and more than capable of producing the results you need, both in game and out.

But was it a breeze to build? There's always some pitfall, right? Some unexpected, frustrating anomaly to deal with on these rigs. Surely the H700i isn't exempt from Murphy's law? Actually, there wasn't much in the way of problems this time around.

Admittedly, it did take us longer than usual to put this beauty together, and that came down to two hiccups. First and foremost: the panels. Both the front and top fascias have the traditional plastic locking pins that secure chassis panels in place, which is fine. However, the ones in the H700i are particularly tricky to pop out. Combine their reluctance to budge with their potential fragility, and you suddenly start panicking about the possibility of snapping them off completely.

The second problem came from that radiator bracket we mentioned earlier. The first time we installed the AIO, we forgot to attach and thread through the fan cables on the Noctua AF-14s we were using. The second time, we noticed the thumb screws wouldn't secure it, because the fans themselves came into contact with the RGB LED light strip NZXT had embedded in the top. The answer in the end was to remove the LED strip entirely (simply two Phillips head screws), then unplug it from NZXT's integrated CAM fan controller.

As for performance? Well, we'll leave the majority of the metrics to speak for themselves, so far as computational performance goes (but more on that in our full Zen+ review on page 74). One thing we haven't really spoken about is the graphics. We decided to go with an MSI GTX 1080 8G for this rig. You can pick one of these up right now for around $600 (probably less by the time you're reading this), and it's fast becoming our 1440p GPU of choice (again).

Is there anything we'd change? That depends on the circumstances. Other than adding better cable combs, the only thing that really stands out is the storage configuration. If this were solely a gaming machine, a 500GB SSD up front and a 2TB hard drive in the back would have made much more sense, and save us a fair bit ditching the Intel drive, too. Otherwise, it's a pretty slick setup for a high-end enthusiast.

BENCHMARKS

	ZERO-POINT	
Cinebench R15 Multi (Index)	2,178	1,832 (-16%)
CrystalDisk QD32 Sequential Read (MB/s)	3,136	2,715 (-13%)
CrystalDisk QD32 Sequential Write (MB/s)	2,126	2,283 (7%)
Rise of the Tomb Raider (Avg fps)	68	48 (-29%)
Total War: Warhammer II (Avg fps)	42	31 (-26%)
Tom Clancy's Ghost Recon Wildlands (Avg fps)	38	30 (-21%)
3DMark: Fire Strike (Index)	6,988	5,456 (-22%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our zero-point consists of an Intel Core i9-7900X, 32GB G.Skill Ripjaws V Series DDR4 @ 3,000, an MSI GeForce GTX 1080 Ti Gaming X, and a 512GB Samsung 960 Pro PCIe SSD. All tests performed at 4K at the highest graphical profile.

CRAFT, BUILD, FIGHT AND EXPLORE IN THIS BRAND NEW MINECRAFT SPECIAL!

Packed full of build guides, gameplay advice and tips from your favourite YouTubers, The Ultimate Guide to Minecraft is an essential companion for any Minecraft fan

132 PAGES OF BRAND-NEW MINECRAFT FEATURES
VOL 20 AUGUST 2017
GAMES MASTER PRESENTS
100% UNOFFICIAL!
THE ULTIMATE GUIDE TO
MINECRAFT
PS4 • XBOX ONE • SWITCH • PC • PS3 • 360 • P
EMBARK ON A VOYAGE OF DISCOVERY
FANTASTIC CREATIVE MODS
10 TOP TOOLS FOR BUILDERS
EXCAVATE AN UNDERGROUND LAIR
BUILD A FLYING MACHINE
TRANSFORM A DESERT TEMPLE
MASH-UP MAGIC
GRAB YOUR FRIENDS, IT'S ADVENTURE TIME!
Explore the Land of Ooo with Finn, Jake and the rest of your favourites
FEATURING
WOODLAND MANSIONS
Dare you face the monsters inside?
EXPERT BUILD ADVICE
YouTube star's top tips revealed
HOW TO STAY ALIVE
Don't fall victim to hostile mobs
Future

ON SALE NOW

FUTURE

Ordering is easy. Go online at:

www.myfavouritemagazines.com

Or get it from selected supermarkets, newsagents or bookstores

REVIEWS

TESTED. REVIEWED. VERDICTIZED.

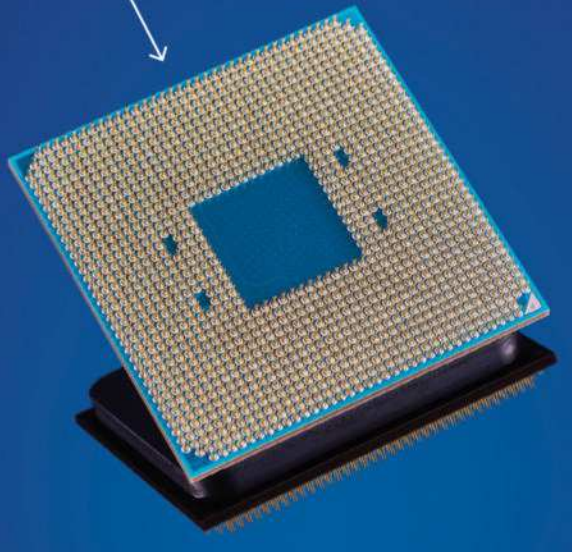

INSIDE

74 AMD Ryzen 7 2700X
76 Asus ROG Zephyrus M GM501
78 MSI X470 Gaming M7 AC
80 Intel Hades Canyon NUC8i7HVK
83 Intel SSD DC P4510 2TB
84 Viewsonic VP3881
87 SteelSeries Arctis Pro + GameDAC
89 Roccat Kone AIMO
90 Final Fantasy XV Windows Edition
91 The Council Episode One: The Mad Ones
92 Lab Notes

AMD Ryzen 7 2700X

It's semi-shrunk, but is it truly second gen?

AMD'S RYZEN was a very long time coming. So much so, some industry observers worried whether AMD would survive long enough to bring it to market. But AMD did, and here we are just over a year later. In the meantime, the original Ryzen CPU family inserted a flying nocturnal insectivore up the metaphorical nightdress of the PC processor market. And now we have AMD's second-generation Ryzen processor in hand. AMD is back, and it's all for the better, even if you don't buy an AMD CPU.

It's arguably a stretch to characterize the Ryzen 7 2700X chip as truly second gen. Architecturally, the narrative involves Zen+ while previous Ryzens were Zen. But there are no changes to the physical design, no new features, no additional cores, no added PCI Express lanes, no increased cache memory, nothing. This chip is all about changes to the production process, plus a few firmware tweaks.

The big upgrade, then, is the shift from manufacturing partner Global Foundries' existing 14nm technology to its latest 12nm node. At which point we partially enter the realm of difficult-to-confirm claims, short of pulling apart the processor, and putting it under an electron microscope, at least. GloFo says tweaks to its 14nm node make for a partial optical shrink, and in turn a 15 percent increase in feature density.

Arguably more important and something we can assess, albeit indirectly, is GloFo's claim of a 10 percent performance improvement with 12nm over the old 14nm process. On paper, AMD says that translates into an overall uplift of around 250MHz, to 4.35GHz, with single-threaded workloads, and all-core overclocks of 4.2GHz. More specifically regarding this eight-core, 16-thread 2700X CPU model, that means a base clock of 3.7GHz, and a maximum Turbo of 4.3GHz.

CLAIMS TO FAME

That compares with 3.6GHz and 4.0GHz respectively for the fastest equivalent chip from the first-generation Ryzen family, the 1800X. For the record, there will be no 2800X processor. This is as good as it gets for now, although AMD has not ruled out a 2800X model in future. But how good, exactly, is the new 2700X? Actually, it's more or less exactly as good as the claims we've just laid out.

Take the 2700X, and slap it into exactly the same PC that was just running an older 1800X. Same motherboard, same memory, same cooling. Then run the numbers. The 1800X scores 1,640 points in Cinebench. Multiply that by the 10 percent performance improvement the new process supposedly delivers, and the result is 1,804 points. The reality? 1,814. That applies pretty much across the board. At least, it does where CPU performance is the deciding factor. The *Total War* game test is clearly bottlenecked by the GPU, not processor performance.

Indeed, this 10 percent paradigm even applies to overclocking. The 1800X fitted to our test rig maxes out at 3.9GHz on all cores. Add 10 percent, and you should see 4.2GHz from the new 2700X, right? Indeed you do. The downside to that predictability is that it means the new 2700X resembles the old 1800X, in not offering any worthwhile overclockability using conventional components and without significant hand-tuning. At 4.2GHz, the 2700X scores 1,888 points in Cinebench. Barely any faster at all than stock clocks.

All of which means this new second-gen Ryzen CPU is not a radical step forward. It's not even a slam dunk in terms of being an obviously better choice than a first-gen Ryzen. The carry-over architecture means the choice is a simple matter of bang for buck. If you could find an older model sufficiently cheaply, it would still be a compelling option. However, the 2700X does build upon all of the original Ryzen CPU's strengths. It delivers slightly more performance for slightly less money. It won't revolutionize your world, but it's unambiguously better than what was already a very good CPU. **–JEREMY LAIRD**

BENCHMARKS

	AMD Ryzen 7 2700X	AMD Ryzen 7 1800X
Cinebench R15 Single/Multi (Index)	**1,814**	1,640
Tech ARP's X264 (Avg fps)	**31.7**	28
Fry Render (m:s)	**2:14**	2:29
AIDA64 Latency (ns)	**68.4**	77.3
AIDA64 Memory Bandwidth (MB/s)	46,506	**46,544**
Power Draw Idle/Load (W)	91/221	**76/199**
Total War: Warhammer II @ 1080p (Avg fps)	**73**	**73**
3DMark Fire Strike Integrated/Dedicated (Index)	**17,855**	16,836
Maximum Overclock (GHz)	**4.2 (1888 CB15)**	3.9 (1678 CB15)
Price to Performance (Index)	**5.50**	5.13

Best scores are in bold. Our test bed consists of an MSI X470 Gaming M7 AC, 16GB (2x 8GB) of Corsair Dominator Platinum DDR4, an Nvidia GeForce GTX 1080, and a Samsung 850 Evo 500GB SSD.
All games were tested at 1080p on the highest graphical profile.

VERDICT

9

KICK ASS!

AMD Ryzen 7 2700X

RYZEN Improves on an already good CPU; more performance for less money; maintains the same socket.

FALLEN Still no overclocking headroom; not super-new design.

$329, www.amd.com

SPECIFICATIONS

Base/Turbo Clock	3.7GHz/4.3GHz
Cores/Threads	8/16
Lithography	12nm
Cache	16MB
Memory Support	DDR4
Memory Channels	Dual-channel
Max PCIe Lanes	24 (16+4+4)
TDP	105W

AMD's second-gen Ryzen slots into the same AM4 socket as before.

Asus ROG Zephyrus M GM501

Six cores of pure gaming sizzle

SIX-CORE INTEL CPU? Check. Nvidia GTX 1070 graphics? Affirmative. IPS panel with 144Hz refresh and G-Sync support? Sir, yes, sir. On paper, Asus's latest ROG gaming laptop, the Zephyrus M, has all the makings of a high-frame-rate portable powerhouse. But like any laptop, what matters is the detailed implementation. Much more than desktops, integration and engineering make all the difference if you're going to squeeze the maximum out of any given component in a laptop.

As it happens, attention to those details is precisely what Asus claims the Zephyrus M delivers. As a for instance, Asus says the thermal design of the chassis compared to its Zephyrus progenitor has been carefully tweaked. The overall thickness is maintained at 17.5–19.9mm, but the volumes have been optimized, and the number of thermal fin outlets has increased to four to improve heat dissipation.

Similarly, the bottom of the chassis has a special heat-dissipating structure. When the screen lid is lifted, a section of the chassis bottom levers open to raise the entire laptop by 9mm, and allow air to flow through the keyboard and maximize circulation. Nice. Speaking of the keyboard, it's fully RGB backlit, and offers 1.7mm of travel. Then there's the 15.6-inch 1080p LCD panel, with claimed 3ms response, and the aforementioned high refresh and G-Sync capabilities, not to mention both a 512GB NVMe system drive and a 1TB SSHD for bulk storage.

What's more, Asus reckons this new portable offers an improved balance between performance and endurance, thanks to a bigger battery and Nvidia Optimus, which switches between the Intel integrated graphics and the Nvidia GPU on the fly. We could continue, but you get the idea. Whether it's features or fine tuning, this thing has it in spades. But does that manage to translate into a great ownership experience?

The performance picture is mostly very impressive. Asus has gone for the Core i7-8750H, one of Intel's new six-core beasts. Under heavy multithreaded CPU loads, and as fitted to this chassis, it'll run at around 3.9GHz for brief periods, before falling back to 3.3–3.4GHz during workloads such as sustained rendering. As for the Nvidia GTX 1070 GPU, we found it typically runs at just over 1,700MHz with occasional dips into the 1,600s. That's true even after a full hour of gaming. Given the desktop GTX 1070 chipset has a 1,733MHz maximum boost frequency, you're clearly not giving up much versus a desktop rig with similar components. The exception involves extreme combined CPU and GPU loads. That can see the CPU clocking back to 2.7GHz.

However, for gaming it's unusual to see both GPU and CPU fully loaded at the same time. So, Asus has done a good job of managing the thermals, even if the area just above the keyboard does get very hot, and the fans kick out some serious noise under heavy loads. It was ever thus with high-performance gaming laptops.

Anyway, physically this is a mostly pleasing portable, thanks to the high-quality chassis that uses plenty of metal, and a very solid-feeling keyboard. The downside involves the IPS LCD panel. The 144Hz refresh and G-Sync plus GTX 1070 graphics certainly make for a slick gaming experience. That's especially true thanks to the 1080p native resolution. It's a good match for the GPU power on offer if very high frame rates are the aim. But the panel itself isn't particularly punchy, and it's surrounded by an oversized bezel. This system is bulky for a 15-incher.

Battery life is another weak area. Despite the ability to switch to integrated graphics, and the larger 55Whr battery, you're looking at approximately 2.5–3 hours of video playback in practice. Long battery life isn't a core competence for this kind of portable, but given the emphasis Asus has put on improving just that, it's disappointing. **–JEREMY LAIRD**

VERDICT 8

Asus ROG Zephyrus M GM501

+ GAME ON Great feature set; super all-around performance; cooling gets serious attention to detail.

– GAME OVER Oversized bezel; merely adequate IPS panel quality; expensive.

$2,699, www.asus.com

BENCHMARKS

	ZERO-POINT	
Cinebench R15 Multi (Index)	743	1,143 (54%)
CrystalDisk QD32 Sequential Read (MB/s)	563	3,327 (491%)
CrystalDisk QD32 Sequential Write (MB/s)	133	1,661 (1,149%)
3DMark: Fire Strike (Index)	13,202	14,495 (10%)
Rise of the Tomb Raider (fps)	99	99 (0%)
Total War: Warhammer II (fps)	62	57 (-8%)
Tom Clancy's Ghost Recon: Wildlands (Avg fps)	48	45 (-6%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our laptop zero-point is the Asus GL502VS-DS71, with an Intel Core i7-7700 HQ, an 8GB GTX 1070, and 16GB of DDR4-2400. All game tests are performed at 1080p at the highest graphical profile.

SPECIFICATIONS

Processor	Intel Core i7-8750H
Graphics	Nvidia GeForce GTX 1070 8GB
RAM	32GB
Screen	15.6-inch 1080p IPS 144Hz
Primary Storage	512GB NVMe SSD
Secondary Storage	1TB SSHD
Keyboard	RGB backlit, 1.7mm travel
Battery	55Whr
PSU	230W adapter
Total Weight	5.4lb

Intel's new six-core mobile CPU combines with the familiar Nvidia GTX 1070 CPU.

MSI's X470 board offers a lot for both first and second-generation Ryzens.

MSI X470 Gaming M7 AC

With a great architecture comes great responsibility

IT'S FAIR TO SAY that the first generation of Ryzen motherboards had its problems at launch. While most of these have since been ironed out, a second-generation chipset brings with it a certain level of reassurance. You just know that those little nips and tucks will make for a better user experience. Indeed, this is exactly what we found when using this motherboard—everything simply worked first time. There's more to this new MSI mobo than the fact it works, though.

It's worth noting that there isn't a lot that's obviously new with the X470 chipset when compared to the older X370. In fact, looking down the spec sheet, there's only one standout inclusion, and that's support for AMD StoreMI storage acceleration—which is essentially a smart caching solution that combines the capacity of your hard drive with the performance of your SSD. As ever with such things, serious performance junkies will get more from a large, fast NVMe SSD, but it's good to see that AMD is thinking outside of its traditional box.

MSI has augmented this base chipset with a plethora of features that elevate it above the norm. You'll find support for RAM speeds up to DDR4-3600, twin M.2 slots, high-end Nihimic 3 audio, Killer Ethernet E2500, Intel Wi-Fi/Bluetooth module, support for dual graphics cards, and even a smattering of RGB lighting. That lighting (termed Mystic Light) is fairly subtle, and if you're determined to make your build shine, there are brighter options out there (the Gigabyte X470 Aorus Gaming 7 lights up like a Christmas tree). Lighting aside, it's a fairly sober design, with the black and silver styling working well in our view.

CASTING SLOTS

The black and silver cooling plate that covers the southbridge extends over the two M.2 slots, providing cooling and protection for your SSDs. The first of these uses the CPU's PCIe 3.0 x4 connection, while the second is routed through the motherboard chipset, and is actually a PCIe 2.0 x4 connection. It's worth noting that you'll lose the third SATA port if you attach a SATA M.2 drive to this connector, or the last PCIe slot if you connect a PCI M.2 drive. And no, you can't build a RAID with these slots, in case you're wondering.

There's a clear confidence on show here that lets you know that you're getting a mature motherboard that will give you the most from the new Zen+ chips. This board is compatible with the original Ryzen family as well, and likewise X370 motherboards can be flashed to support the newer Ryzen 2000-series processors, too. In fact, the new X470 motherboards aren't a direct replacement for the existing X370 motherboards, and both will be on sale going forward.

When it comes to performance, we benchmarked this board with the original Ryzen 7 1800X, so we could easily compare it to MSI's X370 XPower Gaming Titanium (for the Ryzen 7 2700X results on this board, see page 74). The good news, from the off, is that the MSI X470 Gaming M7 AC outperforms, or at least matches, the older board in every test, bar one. The difference can be subtle at times—we're looking at 1 percent in some tests—but better it is. The exception is the power draw, where the older board isn't as greedy as this one.

As this is a halo product, it shouldn't be much of a surprise to find that it's a premium "enthusiast" class motherboard. Rolling in at a cool $260, it's not the cheapest way to get your AMD game on. Given that one of the big selling points of Ryzen is its great value for money, the cost is worth bearing in mind. Overall, this is a confident release from MSI. It boasts a strong feature set, great performance out of the box, and a design that doesn't overdo the RGB lighting like some do. It is on the pricier side, but if you can justify the price to yourself, it comes heartily recommended. **–ALAN DEXTER**

BENCHMARKS

	MSI X470 Gaming M7 AC	MSI X370 XPower Gaming Titanium
X265 (Avg fps)	**28.43**	26.98
Cinebench R15 Multi (Index)	**1,636**	1,615
Fry Render (Seconds)	**161**	**161**
AIDA64 Memory Latency (ns)	**78**	93
Power Draw Idle (Watts)	71	**61**
Power Draw Load (Watts)	179	**166**
CrystalDisk Sequential Read (MB/s)	**553.3**	551.4
CrystalDisk Sequential Write (MB/s)	**535.3**	494.6
Far Cry Primal (Avg fps)	**43**	42
3DMark: Fire Strike Extreme (Index)	**9,934**	9,706
Maximum OC Achieved (GHz)	**4.0**	**4.0**

Best scores are in bold. Our test bed consists of a Ryzen 7 1800X, 16GB (2x 8GB) of Corsair Dominator Platinum DDR4-3200, an Nvidia GeForce GTX 1080, and a 500GB Samsung 850 Evo SSD. *Far Cry Primal* tested on Ultra settings at 4K.

VERDICT **9**

MSI X470 Gaming M7 AC

MYSTICAL Strong feature set; impressive performance; good overclocking headroom; dual M.2 slots; decent audio.

MUNDANE Lighting not to everyone's taste; on the pricey side.

$260, www.msi.com

SPECIFICATIONS

Chipset/Socket	X470/AM4
Form Factor	ATX
Memory Support	64GB DDR4 @ 3,600MT/s
M.2/U.2 Support	2x M.2
SATA Support	6x SATA 6Gb/s
Max PCIe Support	2x PCIe 3.0 x16 (x8/x8)
Rear I/O	Wireless A/C, 2x USB 2.0, 5x USB 3.1 Type A, 1x USB 3.1 Type C, 1x PS/2, 1x Gigabit Ethernet, 7.1 audio out, optical S/PDIF, Clear CMOS, BIOS Flashback+

Intel Hades Canyon NUC8i7HVK

Vega? We'd like you to meet Kaby

WHEN INTEL ANNOUNCED its partnership with AMD to produce a single CPU package with integrated Vega graphics in late 2017, hell's residents must have wondered about the white stuff falling from the sky. But the Hades Canyon NUC is no joke, and packing this much power into such a tiny space is a little like harnessing the power of the atom. Kaby Lake-G is the bombshell of small form factors, thanks mostly to the Radeon Vega M graphics.

Vega M GH configuration is five times faster than Intel's HD Graphics 630, and it's also more than twice as fast as the Vega 11 graphics in AMD's Ryzen 5 2400G. Yeah. It can actually deliver 1080p at 60fps in most games, at medium quality. That's great news for ultra-compact devices like this NUC, and thin-and-light gaming laptops will put Kaby Lake-G to good use. But for our performance table, we're using our mainstream gaming notebook zero-point, which includes a 7th Gen Intel CPU paired with an Nvidia mobile GTX 1070. That comparison shows that while this NUC has the fastest on-package graphics solution ever created, it's still far behind more powerful graphics solutions—the mobile GTX 1070 is almost double the performance. We could also complain that the Vega M is really more like Polaris M, other than the use of HBM2, but the performance is still good.

The laptop comparison is also fitting when you consider the price. Intel sells the i7-8809G NUC for $910 for the barebones system, so you still need to add RAM and storage. Our test unit with SSDs and 16GB DDR4-2400 costs around $1,500—exactly the same as our Asus GL502VS zero-point notebook, which also includes a G-Sync display, keyboard, touchpad, and battery. Clearly, you need to be committed to the ultra small form factor of this NUC, rather than pure performance and features.

Cost aside, the Hades Canyon NUC is pretty swank, with Intel's lit-up skull icon on the top. You can also change the color of the skull, eyes, and power button—or turn off the lighting, if you want something a bit less conspicuous. The NUC8i7HVK packs a ton of IO and connectivity options, including six 4K-capable outputs (two Thunderbolt 3 USB Type-C ports, two Mini-DP 1.2 ports, and two HDMI 2.0a ports). There are also seven USB ports, dual Gigabit Ethernet, an SDXC reader, analog and digital audio ports, and infrared support. It's the kitchen sink approach to NUCs.

Despite all the hardware stuffed inside the petite box, the NUC stays cool and quiet during use, even while pulling up to 150W of total power under load. Hell, you can even overclock the CPU and graphics another 10–15 percent if you're willing to push things. The kit is also Oculus Ready, and the compact size means the NUC can potentially go places other PCs fear to tread, including your living room.

There's only one small snafu if you're thinking of using the NUC as the basis of an ultimate home theater setup. All the 4K video ports are routed through the Vega M graphics, which lacks VP9 profile 2 used for the latest HDR YouTube and Netflix content, as well as the Protected Audio Video Path (PAVP) required for UHD Blu-ray. But 4K streaming is still possible if you're using a supported codec, such as H.264 or HEVC.

If you like cool gadgets and have the money, it's hard not to like this. There are lots of ways to go bigger, faster, and cheaper, but nothing can match the connectivity and performance on tap. The AMD partnership may not stick around long-term, with rumblings of Intel's Arctic Sound dedicated GPU slated for a 2020 launch. Until then, the Core i7-8809G will likely hold the crown for the fastest integrated graphics solution available. **–JARRED WALTON**

BENCHMARKS

	ZERO-POINT	
Cinebench R15 Multi (Index)	743	861 (16%)
CrystalDisk QD32 Sequential Read (MB/s)	563	1,505 (167%)
CrystalDisk QD32 Sequential Write (MB/s)	133	616.8 (364%)
3DMark: Fire Strike (Index)	13,202	8,927 (-32%)
Rise of the Tomb Raider (fps)	99	53.4 (-46%)
Total War: Warhammer II (fps)	62	30 (-52%)
Tom Clancy's Ghost Recon: Wildlands (Avg fps)	48	27 (-44%)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100%

Our laptop zero-point is the Asus GL502VS-DS71, with an Intel Core i7-7700HQ, an 8GB GeForce GTX 1070, and 16GB of DDR4-2400. All game tests are performed at 1080p at the highest graphical profile.

VERDICT **7**

Intel Hades Canyon NUC8i7HVK

NUCLEAR POWER Fastest integrated graphics; good CPU performance; extremely tiny form factor.

NUCLEAR FALLOUT Expensive; slower than most dedicated GPUs; no VP9 profile 2 or PAVP from Vega.

$1,500, www.intel.com

SPECIFICATIONS

Processor	Intel Core i7-8809G
Graphics	AMD Radeon Vega M GH
RAM	16GB (2x 8GB) DDR4-2400
Primary Storage	250GB Intel M.2 PCIe SSD
TDP	100W
Connectivity	4x USB 3.0, 1x SD card reader, 1x Gigabit Ethernet, 1x Mini DisplayPort, 1x HDMI, 1x Thunderbolt USB Type C
Warranty	Three years

Intel's unholy union of Kaby Lake CPU and Radeon Vega M GPU is impressive.

MAKE YOUR OWN PROJECTS WITH THE RASPBERRY PI

Learn the electronics, computing and coding skills you need to make your own projects with the Raspberry Pi, and let your imagination run wild

Ordering is easy. Go online at:

www.myfavouritemagazines.com

Or get it from selected supermarkets, newsagents or bookstores

Intel SSD DC P4510 2TB

The artist formerly known as SFF-8639

ANOTHER MONTH, another new SSD from Intel. Over the last year or two, we've seen Intel roll out new SSDs built into conventional PCI Express cards and SATA enclosures, released in new formats such as M.2, supporting the latest NVMe control protocol in various guises, and even ditching tried and tested NAND flash memory in favor of the radical new 3D Xpoint tech. That's just on the consumer side. Intel's SSDs for servers are even more varied. And yet there's still more.

The "more" in this case is the new Intel SSD DC P4510 Series. It's available in configurations up to a monster 8TB, but this example rocks in at 2TB. Still a beefy drive by solid-state standards. Indeed, the smallest of the P4510 Series is a full 1TB.

In many regards, the P4510 looks fairly familiar. The basic form factor is 2.5 inches, and there's 3D TLC flash memory under the hood. But look elsewhere, and novelties abound. Strictly speaking, it's aimed at enterprise applications, such as cloud servers, but it's still an intriguing device for mere desktop mortals. It hooks up via the lesser-seen U.2 interface, which once seemed like it was set to replace SATA, what with its combination of PCIe-powered pace and a SATA-like simple cable interface, but that never happened.

Actually, the basic character of U.2 is worth recapping. Formerly known as SFF-8639, U.2 supports four gen-3 PCIe lanes and the NVMe protocol, and thus matches the more familiar M.2. Indeed, adapter cards transforming M.2 slots into U.2 ports are available. The difference is U.2's cabled interface, with the drive end of the cable looking a lot like a SAS drive connection, and the motherboard end resembling a double-stacked SATA port.

Beyond the U.2 shizzle, the P4510's specifics go something like this: The NAND memory is Intel's 512GB 64-layer 3D TLC stuff, and the NVMe support extends to version 1.2. Within Intel's hierarchy of enterprise drives, the P4510 falls into the more value-oriented segment, with lower endurance and write speeds than some of its other offerings. It's still very fast in most regards, however. This 2TB version is rated at 3,200MB/s for reads and 2,000MB/s for writes. Read IOPS weigh in at a hefty 624K, while the write IOPS are a mere 79K.

For the target market, it's the endurance rating that may be most relevant. Intel rates it at 0.7 drive writes per day (DWPD). Intel's more expensive NAND drives are rated at around 3DWPD, while the fancy 3D Xpoint-powered Optane P4800X is good for 30DWPD. For the record, this is a 15mm U.2 drive, and thus around twice the thickness of a 2.5-inch SATA drive.

If that's all the speeds and feeds covered off, how does the P4510 actually perform? Assessing it fully from a server or cloud perspective is beyond our remit, but as a client drive, the numbers are pretty zippy. ATTO reveals peak sequential performance very much in line with Intel's claims, although the peak read speed in the CrystalDisk benchmark is disappointing. Random access looks strong, if rather lopsided. But it's the internal file copy that marks the P4510 as something a little bit special: 32 seconds is about half the time taken by a typical consumer M.2 drive, and thus seriously rapid.

Of course, this 2TB SSD is hardly the most cost-effective way of specifying a desktop rig. Almost certainly there are better options. But it does make for a worthwhile reminder of the potential of the U.2 interface, and if the P4510's particular combination of characteristics and capabilities match what would have to be a rather niche set of requirements on the desktop, it might still just make sense. Even at over $1,200. –JEREMY LAIRD

BENCHMARKS

	Intel DC P4510 2TB	Intel 760P 512GB	Intel 800P 1,182GB	Samsung 960 EVO 250GB
30GB Internal Copy (Sec)	**32**	66	76	89
Atto Sequential Read (MB/s)	**3,158**	3,048	1,453	3,019
Atto Sequential Write (MB/s)	**2,059**	1,619	650	1,494
CrystalDiskMark Sequential Read (MB/s)	1,431	**2,811**	1,372	1,682
CrystalDiskMark Sequential Write (MB/s)	**2,032**	1,593	627	1,475
CrystalDiskMark 4K Read (MB/s)	35	75	**229**	64
CrystalDiskMark 4K Write (MB/s)	207	**218**	176	191

Best scores are in bold. Our test bench consists of an Intel Core i7-8700K, an MSI Z370 Godlike Gaming, an Nvidia GeForce GTX 1080, 16GB of Corsair Dominator Platinum DDR4, and a Samsung 850 Evo 250GB OS SSD.

VERDICT **8**

Intel SSD DC P4510 2TB

PURPLE RAIN Big for an SSD; novel feature set; some very strong performance.

THE RAINBOW CHILDREN Very expensive; limited compatibility with existing motherboards; not aimed at client PCs.

$1,259, www.intel.com

SPECIFICATIONS

Capacity	2TB
Interface	U.2
Control Protocol	NVMe
Controller	Intel
NAND Type	3D TLC
Sequential Read/Write	3,200/2,000MB/s
Read/Write IOPS	624/79K
Warranty	Five years

ViewSonic VP3881

Shaking up the high-end super-wide 21:9 segment

EVER SINCE the first 21:9-aspect monitor rolled out five or so years ago, the super-wide format has been about as polarizing as it gets with flat panels. Some loved it for its cinematic sizzle; others trash-talked the limited vertical resolution and consequent poor productivity proposition.

Each successive generation has upped the ante in terms of vertical resolution. ViewSonic's latest cinematic screen salvo is the VP3881. It's a 38-inch monster with no fewer than 1,600 vertical pixels. Thanks to that 21:9 aspect, the horizontal pixel count clocks in at 3,840. In other words, this is a 21:9 respin of the 3840x2160 standard, also known as 4K or UHD.

However you slice it, the VP3881 packs a reasonable number of vertical pixels and, in turn, productivity chops for viewing docs. With this pixel grid, it's practical to view multiple documents side by side. Depending on your vision and comfort levels, as many as four is viable.

Of course, the VP3881 is not entirely novel. Previously, we sampled the LG 38UC99. In many ways, the 38UC99 and this ViewSonic VP3881 are dead ringers. Both sport the same 38-inch (37.5-inch viewable) diagonal, 3840x1600 native res, and are based on IPS panel tech. What's more, both screens have identical 2,300mm curvature radii, and equally offer native 8-bit color depth, with dithering up to 10-bit support. LG is one of the big players in the production of the panel itself, and we do indeed believe that it's one and the same LG panel found in both monitors.

So it's surprising to note that ViewSonic's implementation is unambiguously superior. For starters, it offers slightly wider color spaces. For instance, LG claims 70 percent of the Adobe RGB gamut, while ViewSonic ups that to just over 75 percent. Even more significant, at first glance, is that ViewSonic is claiming HDR10 support. Dig deeper and you'll find that refers to the ability to decode HDR10 video, rather than fully support HDR10 rendering; a glance at the spec sheet reveals a maximum brightness of 300cd/m^2—nowhere near the thousands of nits needed for true HDR playback.

All that said, the VP3881 is noticeably more vibrant and saturated than the relatively disappointing LG 38UC99. Black levels are also superior, with less evidence of the dreaded IPS glow that blights the LG. This is impressive, given the shared panel technology. ViewSonic has also bequeathed the VP3881 a high-quality enclosure and stand, plus good electronics.

The former includes a stand with wide-angle swivel support and 100mm VESA compatibility, not to mention understated but sleek styling. As for the electronics, it's packed with features. Admittedly, the OSD menu controls on the rear of the enclosure come with a learning curve, but after boning up, you'll find some useful features, including a wide range of color space presets—such as sRGB, EBU, SMPTE-C, Rec. 709, and DICOM-SIM—the ability to toggle input lag modes and multiple levels of pixel response, and much, much more.

As for out-of-the box calibration, it's good to go for all but the most discerning pros. ViewSonic claims a Delta E below two for factory calibration, and all the scales and gradients look great to us, except for a little loss of detail in black tones. It ain't cheap, and the pixel pitch is, ultimately, nothing special, but that's an industry-wide issue for desktop screens. **–JEREMY LAIRD**

VERDICT **9**

ViewSonic VP3881

+ HOME CINEMA Genuinely useful vertical resolution; better image quality than LG's panel.

– PEEP SHOW Not as punchy as a true HDR panel; limited to 60Hz refresh.

$1,130, www.viewsonic.com

SPECIFICATIONS

Panel Size	37.5-inch
Native Resolution	3840x1600
Pixel Density	0.229mm
Panel Type	IPS
Maximum Refresh	60Hz
Response	7ms
Contrast	1,000:1
Display Inputs	2x HDMI 2.0, DisplayPort 1.2
Connectivity	3x USB 3, 1x USB C, 3.5mm audio in and out
VESA Mount	100 x 100mm
Warranty	Three years

ViewSonic's implementation of the 38-inch LG panel is that little bit punchier.

DISCOVER THE ALTERNATIVE TO APPLE AND WINDOWS

Take complete control of your computer by learning how to choose, install and use a version of Linux that's tailored just for you

ON SALE NOW

FUTURE

Ordering is easy. Go online at:

www.myfavouritemagazines.com

Or get it from selected supermarkets, newsagents or bookstores

SteelSeries Arctis Pro + GameDAC

Hi-res audio for gaming and music? DAC's entertainment

HIGH-RESOLUTION AUDIO is PC gaming's favorite "Next Big Thing" at the moment. Just as HD displays brought our eyes out of the dark ages and into a new era of crystal-clear game worlds, hi-res audio wants to bathe your ears in newfound fidelity, making positional audio cues in games easier to pinpoint, and Metallica's *Master of Puppets* sound exactly 3.8 times more badass than the last time you heard it. And it's this promise that SteelSeries embraces with its newest batch of Arctis headsets. They were already among the best in the market, and now a trio of fresh (and expensive) models aims to stretch out that lead on the competition.

There's always room for a little healthy cynicism when these new tech standards present themselves, of course. But what we've seen elsewhere from hi-res headsets, such as Roccat's Khan Pro AIMO, is that even if you're not listening to 24-bit sound, and thus capitalizing on the tech, the sound quality is still wonderful. Like that Roccat headset, the Arctis Pro + GameDAC features a frequency response range of 10Hz all the way up to 40kHz, way beyond the usual 20-20 we see from non-hi-res cans. And that means that even if you don't have a Tidal subscription to deliver lossless songs into your ears at the press of a button, you do hear an improvement in sound quality thanks to that wide range.

Best of all, that improvement is coming from the GameDAC box, a digital-to-analog converter and amplifier that takes processing duties away from your PC and handles them in the middle of the signal path. SteelSeries's GameDAC, which is small and unobtrusive enough to sit on any desk, also doubles up as the control center. Pressing the larger black rubberized dial brings up an OLED options menu, where you can turn DTS Headphone X surround sound on and off, tweak game/chat volume balance, and more. As with previous Arctis models, there's a volume scroll wheel at the back of the left earcup, and an easy-to-find mic mute button above it. However, on the wired version we're testing, there's no game/chat balance scroll wheel on the right earcup, and as veteran Arctis 7 users, we feel that loss. It's inherently more tricky to set it using the GameDAC. Another minor gripe is that since the headphones connect to the GameDAC via micro-USB, you need to hook up the supplied 3.5mm mobile adapter if you want to use the headset (minus hi-res support) with your smartphone. And since recent smartphones ditched the 3.5mm jack, that means attaching yet another adapter, and walking around looking like a meme. Not really SteelSeries's fault, but worth bearing in mind if you're considering this as an all-in-one solution.

JACK OF ALL TRADES?

And it is very tempting to see the new Pro version of the Arctis as just that. It's PS4-compatible as well as PC, though not Xbox. Visually, the beautiful clean design hasn't been trodden on by the inclusion of RGB lighting—it's a subtle ring around each earcup that looks as good as we've seen any RGB product look. As a gaming headset, it's feature-rich and brilliantly tuned with the sometimes divisive Headphone X software-based surround sound, creating rich and spacious soundscapes. As a pair of audiophile headphones, it really does have the warmth and depth to compete with more expensive, non-gaming models, albeit with a chain of adapters in tow if you want to hook it up to a smartphone.

This wired Pro model isn't the cream of the new Arctis crop—that honor belongs to the more expensive wireless version, which also includes a chat mix scroll wheel on the headset, and the ease of use that a wireless setup brings. But if your budget doesn't stretch that extra $50, you won't find a more complete package for 250 bucks. –PHIL IWANIUK

VERDICT 9

SteelSeries Arctis Pro + GameDAC

MAN OF STEEL Subtle RGB lighting; great sound; very comfortable.

STEELY DAN No chat mix on earcup.

$250, www.steelseries.com

SPECIFICATIONS

Driver Type	40mm
Impedance	32 Ohm
Frequency Response	10Hz–40kHz
Design Style	Closed cup
Microphone Type	Noise-canceling bi-directional
Connectivity	USB
Weight	13.1oz
Cord Length	5 feet

techradar.pro

IT INSIGHTS FOR BUSINESS

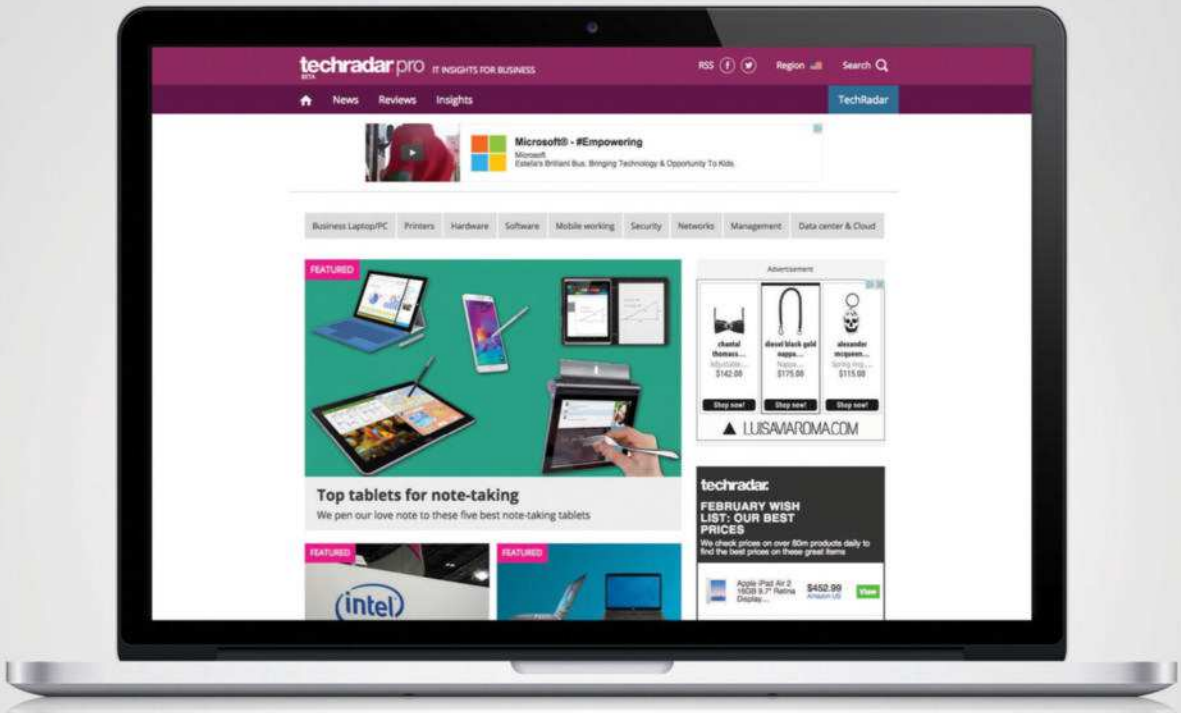

Best practices and solutions

Up-to-the minute news

In-depth reviews

http://www.techradar.com/pro

Roccat Kone AIMO

Eleven years of work produce a thoroughbred

GERMAN FIRM ROCCAT has been making Kones since 2007—the original was the company's first product, and it was the first mouse to include multicolored lighting. So, now we know who to blame.

Roccat is pushing the lighting hard with this model. It "truly travels at the speed of light," gushes the promotional material, something Mr. Einstein might like to have a word about, but there's no getting away from the brightness of the lighting in the Kone AIMO. You're never going to lose this one in a darkened room, and with a 32-bit ARM chip running the show to ensure responsive handling, plus 512Kb memory onboard, there's no reason you'd want to.

The body is thick and curvaceous, pressing firmly up into your palm, and with a deep groove on the left to accommodate your thumb, a ridge at the rear holding it in the perfect position to flick up over the pair of buttons positioned above. The right-hand side is similarly contoured, allowing your ring and pinky fingers to gain significant purchase. Two more buttons nestle behind the wheel, which is broad, and with just enough ridging on its tire—any less would leave it prone to slipping, and we'd like to see a little more.

The wheel tilts and clicks nicely as it goes around, but again, we feel it's holding back and could clunk more firmly into place—that's not to say it's going to be possible to move it by mistake, just that it feels a little limp. For us, it could be positioned a little further forward on the body, too, but most will be pleased with the way it feels. Switches are Roccat's own, with a 50-million-click lifespan, and each button feels responsive, with no sticky movements or shallow actions making you wonder whether your press registered.

The Kone AIMO's big selling point is its lighting. With a lit wheel and two stripes that run along the length of the body, it's almost restrained by modern standards, and produces an attractive glow from between your fingers while in use. It's certainly bright, not losing any of its luminosity as it cycles through colors. AIMO itself is an "intelligent" lighting system that reacts to the way you use the mouse, and is "enriched" (no, us neither) by the apps you use. It does all this in quite subtle ways that are difficult to see when you're concentrating on what you're doing and your hand is blocking most of the light. Other lighting modes, such as simple color shifting or "breathing," are also available from the Roccat software.

SWARM-HEARTED

Whether you get a kick out of intelligent lighting or not, controlling it all through the Swarm software is a breeze. Once it finally finishes updating and installing a new module for the Kone AIMO, Swarm has settings for just about everything, from button assignments to sensor sensitivity (hilariously called the Owl-Eye, it's a 12,000 dpi Pixart 3361 optical sensor that Roccat has tinkered with) and lift-off distance, plus a simple but clever shifting system that defaults to the bottom thumb button, and easily moves to a second programed set of button assignments. Five profiles can be switched between, and there are macros, too, with built-in profiles for some of the best games of 2016. It's a comprehensive package, and we'd have been disappointed with anything less.

The Roccat Kone AIMO is a mouse that does so much right that any criticisms of it are mere niggles. At its price point, the AIMO is taking on the likes of Razer and Logitech, and would be punished for any missteps by losing sales. Roccat has managed to produce something precise, controllable, and customizable (you can even bring back the glory days of the Microsoft Intellimouse by enabling mouse trails, if that's what tickles your fancy), and it's good to see a company produce a solid product that's almost entirely gimmick-free. **–IAN EVENDEN**

VERDICT **9**

Roccat Kone AIMO

ROCKETRY Well designed, well supported, and generally well proportioned.

ROCKABILLY Lighting a little baffling; wheel could be chunkier.

$80, www.roccat.org

SPECIFICATIONS

Sensor	Optical
Sensitivity	12,000 dpi
Sensor Model	Roccat Owl-Eye
Polling Rate	1,000Hz
Programmable Buttons	3
LEDs	16.8 million
Cable Length	6 feet
Weight	4.6oz

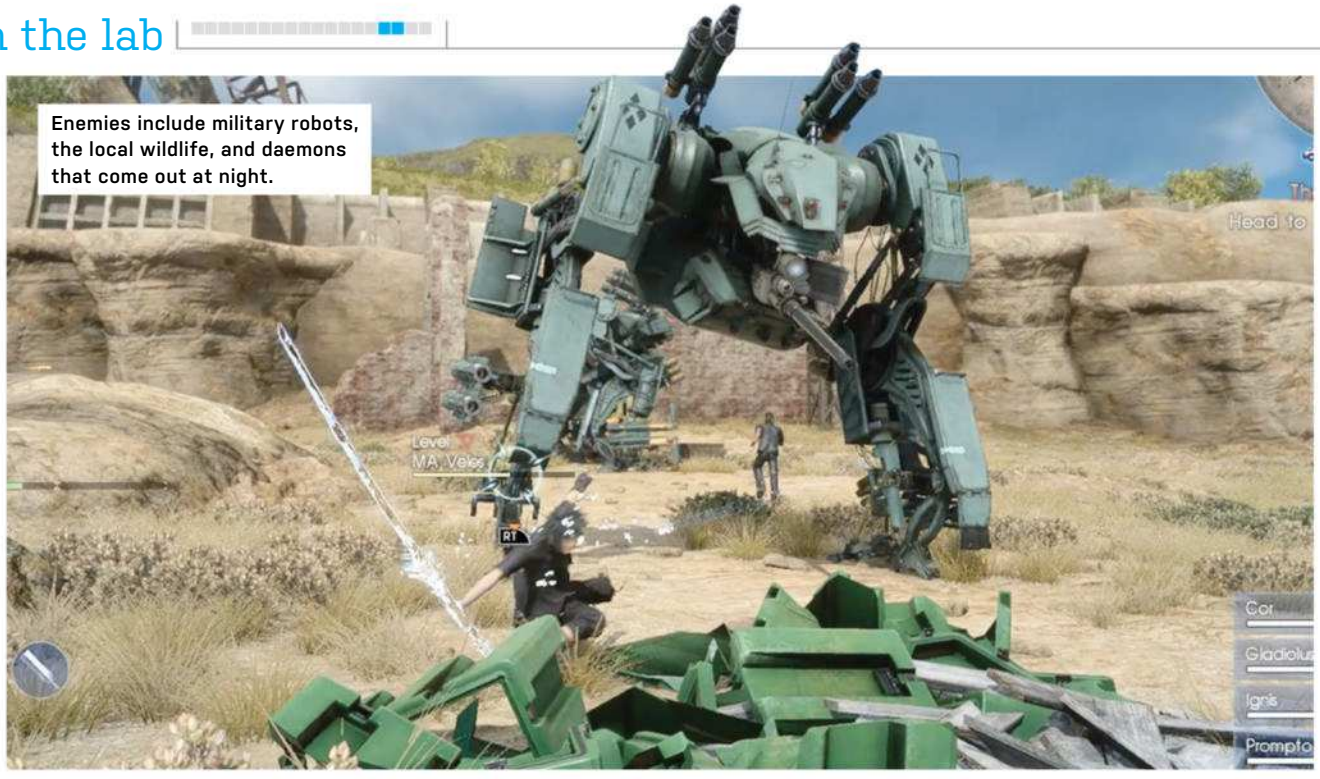

Enemies include military robots, the local wildlife, and daemons that come out at night.

Final Fantasy XV Windows Edition

Unstoppable series requires bottomless hard drive

"A FINAL FANTASY for fans and first timers" is the message you're shown every time you start up the game, and it's a heartening statement to see when approaching a series that has had a troubled relationship with the PC in the past, with some pretty terrible ports. Crimes have included mobile versions of games uprezzed to 4K, while games made for PS3 that should have run at 1080p were locked at 720p.

All that has changed, however, because Square's never-ending series has finally embraced the PC and everything it can do. It's a game that can look stunning, and which can bring even the beefiest rig to its knees with a few tweaks to the settings.

First of all, though, there's the small matter of storage. You might need a new SSD, because with the optional high-res texture pack installed (face it, you're going to download it), *FFXV* takes up 152GB. That's twice what *Shadow of War,* with its high-res textures, occupies—a ridiculous amount of space, but also a hint of things to come. Games are only going to get bigger, and when they look like *FFXV,* we're not too bothered by the groaning of our Internet connection.

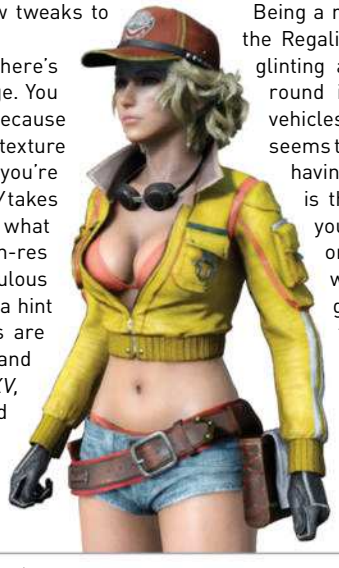

The game is simplified from the convoluted RPGs that came before in the series. There's more focus on action here, and despite the bewildering number of sub-menus, selection screens, and choices, it really is a streamlined game, framed as a road trip taken by four friends—all muscular young men, with fashionable haircuts, and a liking for wearing black—one of whom is a prince on his way to his wedding. Of course, it doesn't work out that way; there's an invasion, people die, and you're soon thrust into a quest to retrieve ancient weapons, and kill a lot of daemons and robots.

Being a road trip, there's a car, and the Regalia is a beauty, shining and glinting as you swing the camera round it. The world is built for vehicles, but almost no one else seems to drive one, so the feeling of having the open road to yourself is thrilling, whether you drive yourself between locations or let your buddy take the wheel. Near the end of the game, you can upgrade the car so it flies, but for most of the game, you're confined to the roads. This is no bad thing, as setting a destination on the map and being driven to the next quest marker or point of interest means you get to listen to the pals chat and bicker, and some great work has gone into making them sound like genuine friends.

Being out on the road is glorious, so it's a shame that the second half of the game becomes a more linear experience involving a train. It's hard to get too upset, however, when the whole thing is backed by such a magnificent score, and the highly tweakable graphics look so good. *Final Fantasy XV* is a console port done right, and despite being a more streamlined game than we're used to from the series, it still deals with subject matter that could rightly be described as "epic," while placing the relationship between the four men front and center. It's a journey worth taking, as long as you've got the storage space. **–IAN EVENDEN**

VERDICT 8

Final Fantasy XV Windows Edition

FANTASY Gorgeous looks and sound; top-quality script; fantastic character relationships.

REALITY Huge storage demands; 4K 60fps probably out of reach; second half rather linear.

RECOMMENDED SPECS i7-3770/FX-8350; 16GB; GTX 1060 6GB/RX 480.

$50, www.finalfantasyxv.com, ESRB: T

None of the characters is all they seem....

The opulent environs are intriguing.

2 Blunders left
ESCAPE THIS ALTERCATION
What? Don't tell me you're going to defend these harlots!?
Play along with him
Let's calm things down.
Keep quiet
There's nothing better than a wise word.
Will you intervene as this brute is rude to a lady?

Your mother: not looking her best, but you love her anyway.

The Council Episode One: The Mad Ones

Why you should never take your mother to a party

BEING INVITED to a lavish party on a private island is something that happens to us only rarely. So when it coincides with meeting a buxom duchess, a creepy cardinal, various historical figures from the 18th century, and a lord who may or may not be immortal, the only thing that can possibly spoil it is your own mother.

Given that you start the game having been tied to her by a mad German, it's clear that you and your mom have adventures together. They're probably something to do with the secret society she runs. But now she's gone missing, and you've come to the private island party in an attempt to find her.

The Council is a narrative adventure game, a third-person point-and-clicker with way more dialogue than violence, and a neat way with social manipulation. It's also episodic, this opening installment taking around three hours to play through if you don't rush, and ending on something of a cliffhanger.

As a member of a secret society, you have a number of skills—assigned by points that you gain as you level up—which can be put to use to interrogate, manipulate, and generally get information out of the partygoers (and the masked servants who make the setting feel so much like something that fell out of *Dishonored*).

The number of skills and the presence of three save slots suggest you're expected to replay the episode, and there are certainly many moments that pass you by if you don't have the requisite number of points invested in the right places. Going through the game with a logical, scientific character won't trigger too many different events from playing as a thief or a charismatic charlatan, but the dialogue options, the way characters react, and whether you can access certain areas differs. Filling up all three saves, then picking one to go on to the next episode would be a difficult choice.

Skills can be topped up by potions you find lying around or in locked boxes, either refilling the Effort Points you spend to use your skills, or making the next use free. Beware using too many top-ups at a time, however, as they can make you intoxicated and remove the countdown bar from time-limited decisions, causing panic. The dialogue is the heart of the game, and many skills affect your reading of another person, so you can say the right thing—converting someone to your point of view through clever speech is the key to several confrontations, and while failure won't end the game, you get fewer experience points, and therefore fewer new skills as a result.

Wandering around, your attention is drawn by glimmers of light to things that can be stolen, read, or generally messed about with. It's straightforward, and is so much better than playing "hunt the interactive pixel" during the golden age of point-and-click adventures. Graphically, with its wood-paneled hallways, dim libraries, and almost comically taut, wrinkled faces that are begging for a bit more antialiasing, the game isn't pushing any boundaries. Its strengths are in its writing, and in the way you're forced to think about each encounter. It's a strong first showing for an episodic series, and now it's up to the developers to keep up the quality. **–IAN EVENDEN**

VERDICT 7

The Council Episode One: The Mad Ones

PARTY TIME Heaps of dialogue; intriguing setting; utterly raddled faces.

PARTY'S OVER Occasional decisions feel like guesswork; over far too soon.

RECOMMENDED SPECS i7-3770/FX-8370; 8GB; GeForce GTX 960/Radeon R9 380.

$30, www.thecouncil-game.com, ESRB: M

LAB NOTES

BO MOORE, TECHNOLOGY EDITOR

Smaller is Better

Why I like 14-inchers

MSI's new GS63 Stealth Thin looks to be the first real gaming answer to the MacBook Pro.

I'VE USED A LOT OF LAPTOPS over the years. From personal machines to gaming systems I've reviewed for these pages and others, I have experienced a huge breadth of portable personal computers. And right now is one of the most exciting times in the laptop space in recent memory.

Why am I so excited? Because I love a laptop that's actually portable, and companies are finally realizing that just shaving a few millimeters off a laptop's thickness doesn't mean much if you still can't fit it in your bag.

I found this painfully true when Nvidia's Max-Q philosophy hit the market around this time last year. Suddenly, laptops were light, thin, and powerful, but their 15-inch bodies made it difficult to carry them in the sorts of bags designed for the Ultrabook form factor. Only one company at the time—Gigabyte, with its Aero 15X—offered a gaming laptop I could toss in my backpack without the corners awkwardly sticking out.

This year, things are different. With the Coffee Lake refresh to most companies' laptop offerings, at least one has followed Gigabyte's lead with a thin-bezel design that similarly fits a 15.6-inch screen in a 14-inch laptop body. MSI is the second out the gate, with its GS65 Stealth. Sleek design, powerful internals, a beautiful, high refresh rate screen, and most importantly, that coveted 14-inch body.

An inch might not seem like much, but it makes a world of difference when it comes to portability. These 14-inchers finally feel like the gaming answer to the ever-popular MacBook Pro.

ZAK STOREY

Deputy Editor

We messed up. We didn't spot an error going to print. I'm saying "we," but if I'm honest, this one rests solely on my shoulders.

Our laptop reviews over the last two issues have gone a little awry: Two figures in the benchmark tables are in the wrong order. Unfortunately, *Rise of the Tomb Raider* didn't score 13,351 avg fps at 1080p—it seems the 3DMark and *Rise of the Tomb Raider* scores ended up being swapped over somewhere along the line. I should have spotted it, because the benchmarks in *Maximum PC* are my responsibility. It's definitely a proofing error on my part, and I apologize.

I'll take this chance to give you an insight into the review process. Typically, I speak to the PR to get an item in, arrange a photography date, commission one of our writers to review it, give them a list of criteria they need to adhere to (specs, benchmarks, word count, and so on), and ship them the item. The copy comes in to me, usually on time. I then proof it, and send it to our designer, Fraser, who lays it out on the page, after which it goes to our sub-editor, Katharine, who proofs it again. At this point, it's ready to go to the big cheese, aka executive editor, Alan Dexter. He checks it over, I check it again, and it's then converted into a high-res PDF, at which point Katharine checks it one last time, before Fraser marks it "ready for print."

Editors' Picks:
Digital Discoveries

Jarred Walton, senior editor, and Alan Dexter, executive editor, share their latest hardware loves

COOLER MASTER ML240L RGB

Some things you don't need in your PC, but damn, do they make life more convenient. That goes double if you like to tinker with your PC. A CPU all-in-one liquid cooler is at the top of my list for every system I build, because I know I'm going to be swapping CPUs.

Mounting or removing a heatsink can be a hassle with large air coolers, and smaller air coolers aren't sufficient for the fastest CPUs—especially with overclocking. Modest AIO coolers don't cost much and tend to perform better, and while installation is more complex, it's the long-term aesthetic appeal and easier access to components that I appreciate. I've encountered too many PCs where it's difficult to access the PCIe latch for the top graphics card due to the size of the cooler. That's a deal breaker.

But which is best? The 280mm and 360mm models provide maximum cooling performance, but Cooler Master's ML240L fits in smaller cases, and still performs well. It's not the quietest of the bunch, and it uses your motherboard's fan connectors instead of a dedicated controller, but I prefer that, as it's one less piece of software to run. What seals the deal is the price—it's about the same as the best air coolers.

$70, www.coolermaster.com

RASPBERRY PI 3 MODEL B+

I ordered the original Pi the day it went on sale. And I loved it. Or rather, I loved the idea of it. It was a bit too sluggish, but when it was released, it was all about its potential. Over the years, I've bought several iterations of the little machine, but with this latest release, I think my reactionary purchases have reached a natural end. Not because there's anything wrong with it, but rather the opposite: Its feature set and performance is now at a point where I'm happy to stick with it.

It's not going to challenge my array of laptops and stupidly powerful desktop machines, but as a tiny machine that I can tinker with without fear, it does everything I want. The kids can go on it and do as much harm as they can manage, and it's only ever a fresh install of the microSD away from being up and running again. It's great for programing projects, old game emulation, browsing, and a myriad other little projects. I'll admit I've never got round to any of the hardware builds I've been collating; they're still an option if I ever get any free time again.

Yup, I'm convinced this will be the last RPi I'll need to buy. It's brilliant.

Who am I kidding? I'm obviously going to buy the next one, whatever that may be.

$35, www.raspberrypi.org

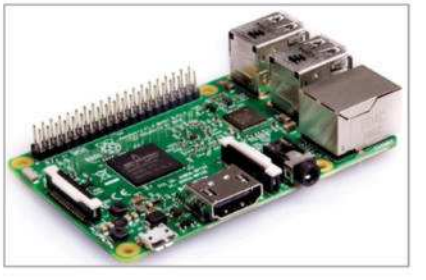

GEEK TESTED & APPROVED

Griffin Survivor Strong for S9+

IF YOU'VE BEEN FOLLOWING my musings, you'll know I'm very much into downhill mountain biking. Getting out into the sticks, and smashing some ludicrous speeds down mountains is possibly the best thing I've ever experienced.

That said, back in 2016, I suffered a very nasty crash. I was on a bike holiday in Scotland, UK, bombing downhill, doing about 20–25mph on a wet, dedicated, slate-covered, single-track trail. I'd sent my friends down ahead of me, as I was the more confident rider, and wanted to "enjoy the moment." Unfortunately, I hit a jump wrong, went over the handlebars, and slid sideways straight into the loose slate, which managed to shred its way through my compression shirt, internally tearing my tricep in half, and lacerating my arm to the point where I was genuinely convinced I'd ruptured my radial artery, there was that much blood.

I got up, screamed a bit, and collected myself. Grabbing my coat, I turned it into a tourniquet, and frantically searched for my phone. That Samsung Galaxy S7 Edge saved my life. Calling the Forestry Commission, they came by pickup truck, bandaged me up, threw me in the back, and punched it straight down to the local hospital.

The thing is, my S7 Edge was perfectly fine. Although it was in my backpack, it was entombed in one of Griffin's Survivor cases. The Survivor series are strong, well-designed cases, designed to take a battering, and I absolutely wouldn't be without one in my life. **–ZS**

$30, www.griffintechnology.com

LETTERS

WE TACKLE TOUGH READER QUESTIONS ON...

> Overclocking Ennui
> Paranoid DNS
> Drive Limits

Why Overclock?

I'm trying to understand something as a non-gamer—well, at least not a gamer in today's world. (Still rocking my old Coleco ADAM, and old 386 and 486 games on PC.) My question is: Why does everyone look at overclocking their CPUs? Why not just buy a processor that's up to the speed that it's rated for? I just don't understand overclocking. What's the deal?

OK, I realize that was more than two questions, but I do have one more: What are "The Builds" all about at the back of the magazine? I see them every month, but I don't understand their story.

–Don Malinowski

EXECUTIVE EDITOR ALAN DEXTER RESPONDS: We'll let you have multiple questions here, but only this once. Never again. And no one else get any ideas on the back of this; we're just feeling magnanimous. To the first part: Why do people overclock? Because they can. Initially, it was a good way of getting better performance out of cheaper chips, and to many degrees that's still the case—why pay more for a high-end chip when you can hit the same performance from an overclocked budget CPU? The theory goes that by saving yourself some money on the processor front, you can spend more where it'll have a bigger impact—on the graphics card, for instance.

There's been a shift here recently, though, as Intel's unlocked chips are considered "enthusiast" parts, which means that it charges a premium for them. Modern chips do have plenty of overclocking headroom, though, which means that they will operate at much higher frequencies than their stock performance. If you need the very best performance around, then overclocking can be the only way to manage it. Having said all that, do you have to overlock? Absolutely not, and the vast majority of users don't bother, but if you fancy the challenge, it can be rewarding, and certainly scratches the tinkerer's itch that many of us have.

As to your second question: What are "The Builds" about? They're essentially a snapshot of the state of the hardware industry at the time they are written. We have three builds each issue, largely defined by their price. There's a budget build, a mainstream system, and a high-end turbo build. The pricing for these systems isn't set in stone (which is a good thing, given the recent volatility in the graphics and RAM markets), but we do try to offer good value for money with each build. If you're looking to build a PC to a particular budget, these systems should give you a rough idea of what is possible at a range of prices.

No Lookup

After a little perusing and benchmarking DNS servers, I started to wonder if I need them at all, or at least all the time—99 percent of the sites I visit, I get there by clicking a bookmark. If I find another page I might want to revisit, I bookmark it. I don't think it's an uncommon approach.

I just substituted Google's IP in a bookmark, and it works fine. Doesn't that skip an extra step, one that has potential for misuse (like every other step along the way)? **–Steve G**

EXECUTIVE EDITOR ALAN DEXTER RESPONDS: You don't need a DNS to go to sites that you know the IP address for, but there's no real benefit to entering the IP address either. Not only does a DNS cover you if a site moves to a different server/IP address, but a DNS also does a good job of giving us humans a more memorable way of referencing websites than trying to recall lots of IP addresses. If I tell you to check out www.pcgamer.com, that's much easier to remember than telling you to go to 185.113.25.60. Think of a DNS as being similar to how you store people's names on your cell, as opposed to just remembering the numbers.

Never-Ending Quest

As a long time subscriber of *Maximum PC*, I was surprised by the article in your April 2018 issue entitled "Industry Wants to Kill Old Online Games." I was not surprised that copyright holders wish to restrict our rights to enjoy old content; I am always opposed to the often-seen overreaching of companies demanding to own everything and charge

↘ submit your questions to: comments@maximumpc.com

us unreasonable amounts to access it. What did surprise me was your including of the game *EverQuest* as an "old online game." While *EverQuest* is indeed one of the oldest online games, it is not abandonware as you imply. The game is still going strong with thousands of players, and expansions released every year. The twenty-fourth expansion was released as recently as December 12, 2017.

Anyone can still play by visiting https://everquest.com and signing up. They can even play free (F2P) with a limited character. In addition, the owners of *EverQuest* have actually allowed fans to run their own version of the game (www.project1999.com). The owners of *EverQuest* have clearly stated that they are OK with Project 1999.

–Bruce Halpern

EXECUTIVE EDITOR ALAN DEXTER RESPONDS: While we'll happily admit to being a little surprised that *EverQuest* is still going strong, we weren't suggesting that it was abandonware at all, simply that the Copyright Office was considering whether to allow the Museum of Art and Digital Entertainment to run old online games. Preserving the likes of *EverQuest* for future generations seems like a reasonable idea, although as you highlight, it would appear that the community is doing a good job of keeping certain old dreams very much alive.

Stealth Mode

In your next review of gaming devices, could you review noiseless gaming mice and keyboards? I know that constant clicking noise is a problem in the office and at home, so this review would really help. **–Alec**

EXECUTIVE EDITOR ALAN DEXTER RESPONDS: We're unashamedly fans of loud keyboards here, although we will keep an eye out for any interesting non-mechanical keyboards, especially as manufacturers have recently started taking an interest in the genre again. However, if any readers have any recommendations on this front, drop us a line, and we'll pass them on.

Speed Test

Tuan's editorial in the April 2018 issue mentions once again the dirty little secret of modern storage: Motherboards that have M.2 slots often share those PCIe lanes with other peripherals and components, making it impossible to run multiple storage devices at full speed.

You brought up this problem in last year's article on PCIe, but I haven't seen it mentioned in any motherboard reviews since then. Perhaps instead of spending part of every review waxing over the board's aesthetics (something we can judge for ourselves from the photograph), you could use those column inches to report how useful the board's multiple M.2 slots actually are for building a RAID 0 out of your lightning-fast Optanes or 960 Pros. In fact, I would suggest that being able to feed at least two NVMe storage devices simultaneously at their maximum throughput ought to be a requirement for a "Kick Ass" rating.

–Stephen Lardieri

DEPUTY EDITOR ZAK STOREY RESPONDS: We actually did do some testing on this way back in our Dream Machine 2017 issue. Plumbing in two 512GB Samsung 960 Pros in RAID 0 provided us with a not-so-stellar maximum throughput of 3,207MB/s sequential reads and 2,146MB/s sequential writes. (For comparison, a single drive manages 3,483MB/s reads and 2,102MB/s writes.) That was using Gigabyte's X299 Gaming 7 and an Intel Core i9-7900X 10-core part. Unfortunately, that was a soft BIOS setup provided by Gigabyte for us, and we didn't actually have Intel's VROC dongle to enable it as a true bootable on-CPU RAID setup, which bypasses the DMI.

The problem stems from the fact that the DMI 3.0 bus connecting the PCH to the CPU (found on everything from Z170 series and up) is limited to a maximum throughput of 3.93GB/s, meaning that a RAID 0 setup utilizing two drives in M.2 slots won't actually gain you anything, and that's unlikely to change until we see DMI 4.0 launch. That's why the majority of add-in card PCIe SSDs only run on a x4 connection, because the DMI becomes the bottleneck at speeds higher than those at which we're already maxing out those PCIe SSDs.

As to why there's so many M.2 slots, given that's the case, this mostly falls down to using PCIe M.2 drives as redundancy or simply as separate drives that you can take advantage of, one at a time.

Slow Drive

Not to travel back in time too much, but I have to ask: What happened on the Digital Storm Bolt X review in the March 2018 issue? The benchmarks against the zero-point show that storage performance was off the mark by almost 70 percent in sequential reads, and nearly as much in writes—those are some large percentages, in my opinion. Yet, the review offers no explanation for these numbers, and the machine still received a "9," which means it should be nearly flawless. Odd. I understand the WD Blue M.2 isn't nearly an EVO 960. That said, I would still hold that the choice of primary drive for the Bolt X is sub-optimal, and that should be reflected in the verdict.

Also, roughly half of the review was about the graphics subsystem. The Bolt X, as reviewed, had an Nvidia GeForce GTX 1080 Ti 11GB—neat, it will play games; I would be shocked as hell if a video card that costs roughly a grand wouldn't. Sort of seems like the review spent a lot of time on the obvious—I mean, this isn't *PC Gamer*—and meanwhile, there is no mention of why two out of the seven benchmarks returned relatively horrible scores (especially for a PC built by Digital Storm, and having a price tag of roughly $3,500). Was this some sort of, "Ah, it's Digital Storm, we'll give them the benefit of the doubt"?

I would love to see honest reviews continue (come back?); it seems as though recently more readers are questioning the verdicts given. Let's keep it fair, and remember: Not everyone only cares about gaming performance.

–Jacob Olson

EXECUTIVE EDITOR ALAN DEXTER RESPONDS: Reviews are very much a product of the time, and given what's been happening in terms of graphics card pricing, it's reasonable that the review spent a good chunk of the words assessing whether the GeForce GTX 1080 Ti made for a quality gaming experience. This actually relates to the other part of your letter—the drive performance. While the throughput from the WD Blue M.2 isn't going to challenge a high-end NVMe M.2 drive, when it comes to gaming, the difference isn't as keenly felt. If this were a machine aimed at video editing, then this would have been a much more important focus, but for gaming, it isn't holding the system back in any way, and so we didn't penalize the system for it.

MAXIMUM PC DIGITAL EDITION

ANYTIME, ANYWHERE!

Delivery of **MAXIMUM PC**, in your choice of
Apple Newsstand, Google Play, or Zinio

From only $1.99 a month

Search for **MAXIMUM PC**
on your device's app store to download now

Available on the App Store

GET THE MAGAZINE ON Google play

zinio™

THE BUILDS

BUDGET

INGREDIENTS

PART			PRICE
Case	BitFenix Nova TG	NEW	$59
PSU	450W Corsair CX450	NEW	$40
Mobo	ASRock AB350M-HDV	NEW	$65
CPU	AMD Ryzen 5 2400G		$163
GPU	N/A		$0
RAM	8GB (2x 4GB) G.Skill RipJaws 4 DDR4-3200		$103
SSD	120GB Kingston A400 2.5-inch SSD	NEW	$37
HDD	Hitachi Deskstar E7K1000 1TB 7,200rpm	NEW	$39
OS	Ubuntu Desktop Linux 16.04 LTS 64-bit		$16

Approximate Price: $522

WE'VE BEEN THREATENING to do this for ages, but we've finally pulled the trigger, and upgraded all of the machines to use newer, more feature-packed cases. When it comes to the budget build, this means a shift from ITX to a full ATX chassis. There's a better range of choice when it comes to ATX, so it's easier to save money—our ASRock motherboard is $20 cheaper than the one we used last month. The move to a full ATX case gives you more room to play with, especially as there's no graphics card here getting in the way. You shouldn't need to add to the default configuration of two front-mounted intake fans and rear exhaust. The ASRock AB350M-HDV is a microATX motherboard, and comes flashed to support 2000-series Ryzen APUs—you don't have to do any CPU juggling to get up and running. With a few savings on the storage front as well, you're looking at a $50 decrease over last month. It's a lot of machine for not much cash.

MID-RANGE

INGREDIENTS

PART			PRICE
Case	Fractal Design Meshify C	NEW	$79
PSU	550W Corsair RM550X		$100
Mobo	MSI Z370 Tomahawk		$150
CPU	Intel Core i5-8400		$179
Cooler	Be Quiet! Dark Rock 3		$75
GPU	Gigabyte GeForce GTX 1060 3GB		$270
RAM	16GB (2x 8GB) Patriot Viper Elite DDR4-2400	NEW	$155
SSD	250GB Samsung 960 Evo M.2 PCIe SSD		$120
HDD	Hitachi Deskstar E7K1000 1TB 7,200rpm	NEW	$39
OS	Windows 10 Home 64-bit OEM		$100

Approximate Price: $1,267

THE NZXT S340 ELITE that has formed the basis of our mid-range build for the last few months is a great case, but we wanted something a little more impactful moving forward, so have elected to give Fractal Design's Meshify C an airing. It's a great case to build in, with enough neat touches to make any build look professional, while maintaining great airflow. There's plenty of room for long graphics cards and CPU tower coolers, and thanks to a particularly spacious rear compartment, you can easily keep your cabling tidy, too. The only other notable change is on the RAM front, where we managed to grab a Patriot kit to save $13 on last month's build. We also saw our graphics card of choice drop another $10, indicating that the crazy mining times may finally be behind us. Making the shift to the Hitachi hard drive saved us a few bucks over last month's WD Blue, too, which overall meant we managed to save $31 compared to last month.

UPGRADE OF THE MONTH

NVIDIA GEFORCE GTX 1080

We've seen graphics card pricing start to return to reasonable levels at the lower end of the stack, while at the top-end, the 1080 Ti has actually got a bit more expensive this month. This makes the non-Ti version of the GeForce GTX 1080 a tempting offer right now, with some manufacturers' cards available for $600 (such as the Gigabyte card pictured). It's still more than the $500 we feel it should cost at this stage in its life cycle, but it's definitely heading in the right direction. If you're rocking a 1440p screen, this is a great card to buy, because it'll handle the highest settings at a silky smooth 60fps. The only potential fly in the ointment is Volta, which should hopefully be heading our way soon, although with no real reason for Nvidia to replace its current lineup, the next-gen cards could be put on hold until AMD manages to force big green's hand.

$600, www.nvidia.com

THE NZXT H700I is everything you could want from a high-end case: packed with features, looks great, and is a joy to build in. The included CAM smart fan control gives complete control over your cooling, too. The only downside is that this is a mid-tower case, which means the EATX motherboard we've been using until now has had to be put to one side, as it's too tight a squeeze for us to be comfortable with. We've grabbed the ASRock X399 Taichi in its place, which is still a large motherboard, but it fits better than the outgoing Asus Prime X399-A.

The new motherboard boasts plenty of high-end features, including an 11-power-phase design to ensure you get the best overclocking potential from your CPU. This is backed up by an eight-layer PCB design to ensure signal stability, which is particularly useful when it comes to RAM support, and is part of the reason we've gone up to DDR4-3200, as this gives the Threadripper CPU a serious boost.

Speaking of Threadripper, the price drop we've recently enjoyed has disappeared, so it's back to $750. We could have dropped to the 1900X, but the extra cores on offer are too good to let go for $80, and it still offers incredible value if you want serious performance. We also saw an increase on the graphics card front, and we were tempted to drop to a non-Ti GeForce GTX 1080, but we've stayed with the faster card for now. Overall, this build rolls in at $350 more than last month, which isn't ideal, but still offers great value for workstation performance.

For more of our component recommendations, visit www.maximumpc.com/best-of-the-best

INGREDIENTS

PART			PRICE
Case	NZXT H700i	NEW	$176
PSU	Corsair RM750X Modular Gold		$110
Mobo	ASRock X399 Taichi	NEW	$349
CPU	AMD Ryzen Threadripper 1920X		$750
Cooler	NZXT Kraken X62	NEW	$156
GPU	EVGA GeForce GTX 1080 Ti SC2 Gaming	NEW	$950
RAM	32GB (4x 8GB) G.Skill Ripjaws V Series DDR4-3200	NEW	$367
SSD	512GB Samsung 960 Pro M.2 NVMe		$325
HDD	2x 2TB Seagate BarraCuda ST2000DM006		$120
OS	Windows 10 Home 64-bit OEM		$100

Approximate Price: $3,403

Maximum PC (ISSN 1522-4279) is published 13 times a year, monthly plus a Spring issue following the March issue, by Future US, Inc., 11 Hanover Square, 14th Floor, New York, NY 10005, USA. Website: www.futureus.com. Periodicals postage paid in San Bruno, CA, and at additional mailing offices. Newsstand distribution is handled by Curtis Circulation Company. Basic subscription rates: one year (13 issues) US: $24; Canada: US$40; Foreign: US$40. Canadian and foreign orders must be prepaid. Canadian price includes postage and GST (GST #R128220688). PMA #40612608. Subscriptions do not include newsstand specials. POSTMASTER: Send changes of address to *Maximum PC*, PO Box 5852, Harlan, IA 51593-1352. Standard Mail enclosure in the following editions: None. Ride-Along enclosure in the following editions: None. Returns: IMEX Global Solutions, PO Box 25542, London, ON N6C 6B2, Canada. Future US, Inc. also publishes *@Gamer*, *MacLife*, *The Official Xbox Magazine*, and *PC Gamer*. Entire contents copyright 2018, Future US, Inc. All rights reserved. Reproduction in whole or in part is prohibited. Future US, Inc. is not affiliated with the companies or products covered in *Maximum PC*. Reproduction on the Internet of the articles and pictures in this magazine is illegal without the prior written consent of *Maximum PC*. Products named in the pages of *Maximum PC* are trademarks of their respective companies. PRODUCED IN THE UNITED STATES OF AMERICA.

ALIENWARE

17

ENGINEERED TO WREAK HAVOC.

The new Alienware 17 is ready for VR, features innovative Tobii eye tracking, boasts a new premium industrial design and has been upgraded with the latest NVIDIA 10-series graphics for a gaming experience so visceral and real, you'll have to decide how deep you're willing to go. **We're Game.**

START GAMING AT ALIENWARE.COM

Alienware is a trademark of Dell Inc. Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of the Intel Corporation in the U.S. and/or other countries.

Intel Inside®.
Extraordinary
Performance
Outside.

We build the world's most advanced PCs.

Experience a new level of performance with an award-winning Digital Storm PC. Built with the latest technology, highest quality components and backed by lifetime support. Visit our website and build your dream PC today.

DIGITALSTORM

LEARN MORE: WWW.DIGITALSTORM.COM

Digital Storm PCs featuring Intel® Core™ i7 processors.

intel CORE i7 inside

The Digital Storm Logo, and 'World's Most Advanced PCs' are trademarks of Digital Storm.Intel, the Intel Logo, Intel Inside, Intel Core, and Core Inside are trademarks of Intel Corporation in the U.S. and/or other countries.